

VOLTERRA–PERRON PRICING OF THE OPTION MANIFOLD

J. VIDAL LLAURADÓ 

ABSTRACT. Public option surfaces can be calibrated without determining the latent risk-neutral $\mathbb{Q}$ -coordinates that move non-public stress claims. This paper treats the option manifold as the public image of a latent Volterra–Perron stress state and formulates recovery as a quotient problem. At date t , the canonical state is

$$\Xi_t^{\mathbb{Q}} = (G_t^{\mathbb{Q}}, \mathfrak{H}_t^{\mathbb{Q}}, \rho_t^{\mathbb{Q}}, e_t^{\mathbb{Q}}, \ell_t^{\mathbb{Q}}, \varpi_t^{\mathbb{Q}}),$$

with a directed Volterra contagion operator, a rough maturity-scaling state, a Perron spectral triple, and a synchronized stress coordinate. A residual-density selector is appended only after recovery. Public calibration does not fix this selected pricing coordinate.

The governing theorem identifies the public-observable quotient. A Volterra–Perron source state is locally recoverable, up to canonical scale, sign, basis, and Perron-normalization gauges, exactly when its source quotient embeds into that public quotient. Stable recovery is positivity of the quotient observability modulus. Without source-family restrictions, public-equivalent density families preserve the full public option manifold while moving latent-stress claims. With a gauge-normalized Volterra source chart, a simple well-conditioned Perron eigenvalue, nondegenerate carrier Gram matrix, maturity-response rank, and spanning index/single-name synchronization responses, the canonical risk-neutral state is locally and stably recovered.

Recovery leaves pricing and hedging completion as separate conditions. Residual price intervals collapse only at public payoff-module completion; otherwise the public-annihilator fiber remains. On recovered state paths, discounted option price functionals satisfy martingale drift restrictions, phase transitions add drift and covariance charges, and selected prices admit native sensitivities to Perron stress, spectral radius, carrier direction, roughness, directed Volterra kernels, and the residual selector. The synchronization relative-value identity decomposes index option value relative to a carrier-replicated single-name book into synchronized-stress premium, residual compression charge, and tradeable basis error. A finite Bachelier–Volterra–Perron benchmark gives a closed option formula, a recovered-state pricing equation, explicit native Greeks, and a finite public hedge residual.

1. INTRODUCTION

Modern option pricing has become increasingly good at fitting the public surface. Dynamic surface models describe the surface itself [1]; rough volatility explains the fractional short-maturity signature in skews and realized volatility paths [2, 3]; affine Volterra and rough Heston models give tractable non-Markovian pricing equations [4, 5, 6]; and neural or surrogate calibration maps interpolate option prices at scales unavailable to early parametric models [7]. These advances improve the public fit. They do not settle which latent risk-neutral coordinates the surface has identified.

The distinction is sharp on an index/single-name option manifold. Single-name surfaces determine marginal pricing laws. They do not determine directed covariance routes, carrier absorption, or synchronized index stress. Index and basket options add joint information, but only through the public response coordinates admitted by the carrier design. In incomplete markets, exact calibration of that public manifold still need not price a claim with nonzero loading on a latent stress direction; unspanned risk appears inside the option manifold itself [8, 9, 10].

Date: May 2026.

2020 Mathematics Subject Classification. 60G22, 91G20, 91G80, 47A75, 47N30.

Key words and phrases. risk-neutral pricing, option manifolds, Volterra operators, Perron stress, latent-state recovery, residual kernels, synchronization risk, incomplete markets.

DOI: 10.2139/ssrn.6744119.

The companion compression papers isolate the same obstruction from the pricing side. Risk-neutral compression identifies finite $\mathbb{Q}$ -compressions of latent Perron stress from public option data [11]. Residual pricing shows that public option prices can fix $C_{\mathcal{H}}^{\mathbb{Q}}(L_X^{\mathbb{Q}})$ while leaving the full latent loading $L_X^{\mathbb{Q}}$ movable along a public-annihilator fiber [12]. The present paper supplies the risk-neutral option-manifold geometry that generates those compressions. The option surface is a measurement of a latent Volterra–Perron state, not the state itself.

At date t , the public object is an option manifold $\mathcal{M}_t^{\text{opt}}$ containing admitted index, sector, carrier, and single-name option prices. The latent source state is

$$\Xi_t^{\mathbb{Q}} = (G_t^{\mathbb{Q}}, \mathfrak{H}_t^{\mathbb{Q}}, \rho_t^{\mathbb{Q}}, e_t^{\mathbb{Q}}, \ell_t^{\mathbb{Q}}, \varpi_t^{\mathbb{Q}}).$$

Here $G_t^{\mathbb{Q}}$ is a directed Volterra contagion operator, $\mathfrak{H}_t^{\mathbb{Q}}$ records rough maturity scaling, $(\rho_t^{\mathbb{Q}}, e_t^{\mathbb{Q}}, \ell_t^{\mathbb{Q}})$ is the Perron spectral triple of the screened synchronization operator, and $\varpi_t^{\mathbb{Q}}$ is the synchronized latent stress coordinate. A selected pricing state $\Theta_t^{\mathbb{Q}} = (\Xi_t^{\mathbb{Q}}, \theta_t)$ appends a residual density selector inside the public-equivalent family. Public prices are generated by the adapted response map

$$\mathfrak{P}_t : \Theta_t^{\mathbb{Q}} \mapsto \mathcal{M}_t^{\text{opt}}.$$

The finite Volterra source chart gives this notation a specified representation: rough kernels produce integrated covariance routes, basket and carrier prices observe finite linear functionals of those routes, and residual selectors are bounded post-recovery pricing rules built from the recovered canonical coordinates.

Calibration is therefore not a single inverse. The public fit maps a state to observed quote coordinates. Source recovery asks whether the gauge quotient of the Volterra–Perron source state embeds into the public-observable quotient. Residual selection chooses a point in a public-equivalent pricing fiber. Tradeable completion projects terminal payoffs onto the admitted instrument span. These four maps may collapse in a complete finite model, but they are distinct on the option manifold considered here.

The central statement is the Volterra–Perron option-manifold equivalence theorem. Public option prices identify exactly the public quotient $\mathfrak{X}_t^{\text{pub}}$. A source family $\mathfrak{S}_t$ is locally recoverable only through the induced map

$$\mathfrak{S}_t / \sim_g \longrightarrow \mathfrak{X}_t^{\text{pub}}.$$

In a differentiable source chart the infinitesimal object is the quotient observability operator

$$\overline{\mathcal{O}}_t^{\mathbb{Q}} = D(\mathcal{R}_t \circ \psi)(u_0) \quad \text{on the gauge quotient.}$$

Recovery holds when this operator has no source kernel and a positive inverse modulus. In the block chart, the condition decomposes into rough maturity, directed-Volterra, carrier, and Perron synchronization blocks. Single-name vanilla surfaces identify marginal laws; they do not identify these joint blocks by themselves.

Pointwise inverses are then assembled through time. Uniform recovery windows give $\mathcal{H}_t$ -adapted recovered processes

$$\widehat{\Xi}_t^{\mathbb{Q}} = \mathcal{I}_t(\mathcal{M}_t^{\text{opt}})$$

and pathwise stability under quote perturbations. The same modulus controls regularized finite panels and gives lower bounds: no reconstruction rule using the same quote data can beat the inverse-modulus scale. Ill-conditioning at carrier-thin, Perron-gap, maturity-rank, or public-completion seams is an information limit, not a defect of a particular numerical inverse.

No-arbitrage enters twice. Static public manifolds must satisfy the convex-order restrictions implied by discounted martingales. After recovery, option prices written as functionals of the canonical state must have discounted local-martingale dynamics. In a finite recovered chart this gives

$$\partial_t \Phi_a + D\Phi_a[b] + \frac{1}{2} \text{Tr}(\Sigma^* D^2 \Phi_a \Sigma) - r_t \Phi_a = 0$$

for each admitted option coordinate. Smooth phase allocations add transition drift and covariance terms because the phase weights themselves move with the recovered $\mathbb{Q}$ -state.

Residual non-identification remains after state recovery. A dynamic public-equivalent density family preserves every admitted public option payoff at date t and across maturities $S > t$ while moving a latent-stress claim whenever that claim has nonzero loading on the public-annihilator fiber. A selector θ_t , measurable with respect to public and recovered coordinates, selects a price inside the residual interval. This gives the dynamic form of the residual-pricing geometry of [12], with no second calibration of the public surface.

The recovered quotient also carries model-native sensitivities and a relative-value identity. Native Greeks differentiate selected prices with respect to Perron stress, spectral radius, carrier direction, roughness, directed Volterra kernels, and the residual selector. The synchronization identity decomposes the value of an index option relative to a carrier-replicated single-name book into synchronized-stress premium, residual compression charge, and basis error. The finite Bachelier–Volterra–Perron benchmark closes these objects in a tractable chart with a closed option formula, recovered-state pricing equation, explicit Greeks, and a finite hedge residual. Its role is the quotient-theory chart; universal smile modelling is outside the claim.

Organization. Section 2 defines the dynamic option manifold and information structure. Section 3 introduces the risk-neutral Volterra–Perron state and its gauge equivalence. Section 4 gives a Volterra–Perron source chart. Section 5 defines risk-neutral market phases and transition surfaces. Section 6 defines the public-observable quotient and proves the public non-identification and public-completion results. Section 7 develops risk-neutral observability, vanilla insufficiency, the minimal option-manifold hierarchy, and finite-quote stability. Section 8 gives the quotient recovery, finite-panel, inverse-modulus, and regularized reconstruction theorems. Section 9 turns pointwise recovery into an adapted recovered state process. Section 10 derives the martingale-consistency drift restrictions for recovered risk-neutral state dynamics. Section 11 gives the canonical finite Bachelier–Volterra–Perron benchmark with its formula, pricing equation, Greeks, and hedge residual. Section 12 verifies the recovery conditions in a reduced Gaussian–Volterra class. Sections 13–15 develop residual kernel selection, native Greeks, public-neutral hedging, and the synchronization relative-value identity. Section 16 gives one equivalent-martingale-measure specialization, and Section 17 records reduced examples.

2. DYNAMIC OPTION MANIFOLDS

Fix a filtered probability space

$$(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{Q})$$

supporting all risk-neutral objects. Conditional-expectation and martingale conventions follow [13]. The public market information is $(\mathcal{H}_t)_{t \geq 0}$, the latent stress information is $(\mathcal{K}_t)_{t \geq 0}$, and

$$\mathcal{H}_t \subseteq \mathcal{K}_t \subseteq \mathcal{F}_t.$$

The tradeable public span $\mathcal{S}_t \subseteq L^2(\mathbb{Q}, \mathcal{H}_t)$ is the closed Hilbert span generated by publicly tradeable instruments available at date t . It may be finite dimensional in applications; the projection statements require only closedness. This span is a public-feature projection space, not the full dynamic self-financing gain space. Section 15 separates recoverability, tradeability, and hedgeability when terminal payoffs and dynamic strategies enter.

Let

$$B_t := \exp\left(\int_0^t r_s ds\right)$$

be the money-market account. Unless an undiscounted payoff is explicitly displayed, all payoff variables in Sections 2–15 are understood in date- t money-market units:

$$\tilde{X}_a := \frac{B_t}{B_{T(a)}} X_a^{\text{cash}}.$$

Thus X_a below denotes $\tilde{X}_a$, and public prices are conditional expectations of discounted payoffs. Section 16 writes the discount factor explicitly again when the underlying cash dynamics are introduced.

Definition 2.1 (Public option manifold). Let $\mathcal{I}_t$ be the date- t list of underliers, strikes, maturities, and option types admitted to the public surface. A public option manifold is the $\mathcal{H}_t$ -measurable element

$$\mathcal{M}_t^{\text{opt}} = (C_t(a))_{a \in \mathcal{I}_t} \in \mathcal{Y}_t,$$

where $C_t(a)$ is the arbitrage-free risk-neutral price of payoff X_a and $\mathcal{Y}_t$ is the corresponding finite- or infinite-dimensional quote space.

Definition 2.2 (Quote-space norm). The quote space $\mathcal{Y}_t$ is equipped with a fixed norm. For a finite panel this may be a weighted Euclidean norm, for example bid-ask or reliability weighted. For a continuum maturity and moneyness manifold it is a weighted C^k , Sobolev, or Banach-supremum norm on the admitted coordinate domain. Dynamic stability statements use a deterministic ambient quote space, or a Banach bundle with uniformly equivalent local norms, over the recovery window.

Definition 2.3 (Observable response coordinates). Write a quote coordinate as

$$a = (u, \tau, m, \kappa, \chi),$$

where u is an admitted index, sector, or single-name underlier, $\tau > 0$ is maturity, m is the moneyness coordinate, κ records call, put, variance, or surface-functional type, and $\chi \in \{\text{B}, \text{N}\}$ records the pricing chart. Here B denotes the Black, or lognormal-forward, chart and N denotes the Bachelier, or normal-forward, chart. In the Black chart one may take $m^{\text{B}} = \log(K/F_{u,t,T})$; in the normal chart one may take $m^{\text{N}} = K - F_{u,t,T}$, or its deterministically rescaled version. For variance and surface-functional coordinates, χ records the convention in which the reported total variance is expressed. For a latent state Θ , the observable response map is

$$\mathcal{R}_t(\Theta)(u, \tau, m, \kappa, \chi) := C_t^\Theta(u, \tau, m, \kappa, \chi).$$

The option manifold $\mathcal{M}_t^{\text{opt}}$ is the realized public value of this response map on the admitted coordinate set. Continuous maturity/moneyness manifolds and finite quoted panels are treated as two observation regimes for the same map.

Assumption 2.4 (Arbitrage-free public manifold). For each t , the model works under a selected equivalent martingale measure $\mathbb{Q}$ on the public payoff span, so the prices $C_t(a)$ are represented by $\mathbb{Q}$ -conditional expectations. Equivalently, the public manifold is a no-arbitrage price system on its admitted public payoff family, and $\mathbb{Q}$ is a selected representative of the equivalent pricing class [14, 15, 16, 17]. Residual selectors introduced below are conditional density changes around this representative that preserve the admitted public prices.

Definition 2.5 (Public no-arbitrage cone). For a finite public payoff panel, write $X = (X_1, \dots, X_M)$ and let

$$\mathcal{A}_t := \left\{ p \in \mathbb{R}^M : h \cdot X \geq 0 \text{ } \mathbb{Q}\text{-a.s. and } h \cdot X > 0 \text{ with positive probability imply } h \cdot p > 0 \right\}$$

be the strict no-arbitrage cone of public prices. Its tangent cone at a price vector $p_0 \in \mathcal{A}_t$ is denoted $T_{p_0}\mathcal{A}_t$.

Proposition 2.6 (No-arbitrage tangent condition). *Let $F : U \rightarrow \mathbb{R}^M$ be a finite-panel option response map whose image lies in the no-arbitrage cone $\mathcal{A}_t$, and let $p_0 = F(u_0)$. Then every source tangent direction that is visible in public quotes satisfies*

$$DF(u_0)h \in T_{p_0}\mathcal{A}_t.$$

If p_0 is in the relative interior of $\mathcal{A}_t$, this imposes no first-order inequality. At the boundary of the cone, it removes source directions whose first-order public response would create a static arbitrage at first order.

Proof. For every sufficiently small $s > 0$ with $u_0 + sh \in U$, the response $F(u_0 + sh)$ lies in $\mathcal{A}_t$. Therefore

$$\frac{F(u_0 + sh) - p_0}{s} \longrightarrow DF(u_0)h$$

is a tangent direction to $\mathcal{A}_t$ at p_0 . If p_0 lies in the relative interior of the cone, the tangent cone is the linear span of $\mathcal{A}_t - p_0$, so no additional first-order inequality is imposed. At a boundary point, the tangent cone is a proper cone and excludes first-order responses that point outside the public no-arbitrage region. $\square$

Definition 2.7 (Conditional convex order). Let μ_{t,T_1} and μ_{t,T_2} be conditional laws, given $\mathcal{H}_t$, of a discounted underlying vector at maturities $T_1 < T_2$. Write

$$\mu_{t,T_1} \preceq_{\text{cx}} \mu_{t,T_2}$$

if, for every convex function φ for which both conditional expectations are finite,

$$\int \varphi(x) \mu_{t,T_1}(dx) \leq \int \varphi(x) \mu_{t,T_2}(dx)$$

pathwise, conditionally on $\mathcal{H}_t$.

Theorem 2.8 (Calendar convex-order restriction). *If the discounted underlying vector is a $\mathbb{Q}$ -martingale and $\mu_{t,T}$ denotes its conditional law given $\mathcal{H}_t$, then, for $t < T_1 < T_2$,*

$$\mu_{t,T_1} \preceq_{\text{cx}} \mu_{t,T_2}.$$

For a scalar underlying this implies the usual calendar inequalities for call prices,

$$C_t(K, T_2) \geq C_t(K, T_1),$$

after the common-forward convention is imposed. For a vector underlying, every linear projection must satisfy the corresponding one-dimensional convex-order restriction.

Proof. For convex φ , Jensen's inequality and the martingale property give

$$\mathbb{E}^{\mathbb{Q}}[\varphi(\tilde{S}_{T_2}) \mid \mathcal{F}_{T_1}] \geq \varphi(\mathbb{E}^{\mathbb{Q}}[\tilde{S}_{T_2} \mid \mathcal{F}_{T_1}]) = \varphi(\tilde{S}_{T_1}).$$

Taking conditional expectation given $\mathcal{H}_t$ gives the convex-order statement. The scalar call form follows from the representation of convex functions by affine functions and positive mixtures of calls. Linear projections of vector martingales are martingales, so the same argument applies to every projection. $\square$

Corollary 2.9 (Dynamic falsifiability before source fitting). *A multi-maturity public option manifold that violates conditional convex order cannot be generated by any equivalent martingale measure, regardless of the Volterra–Perron source family. Conversely, a source-family fit that matches maturity slices separately but violates the convex-order monotonicity across maturities is dynamically inadmissible.*

Proof. The conditional convex-order inequalities of Theorem 2.8 are necessary for every equivalent martingale measure generating the discounted underlier. They are imposed on the public manifold before any Volterra–Perron source family is chosen. A violation therefore rules out the manifold at the martingale level. A fit that matches separate maturity slices but breaks those inequalities produces the same obstruction, now inside the fitted source family. $\square$

Definition 2.10 (Latent loading and public compression). For $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$, define the latent loading at date t by

$$L_{t,T}^{\mathbb{Q}}(X) := \mathbb{E}^{\mathbb{Q}}[X \mid \mathcal{K}_t],$$

and its public compression by

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

The tradeable public projection is $C_{\mathcal{S}_t}^{\mathbb{Q}} X$, the orthogonal projection of X onto $\mathcal{S}_t$ in $L^2(\mathbb{Q})$.

Proposition 2.11 (Compression factorization of public prices). *If an option payoff X_a is priced at date t by conditional expectation under $\mathbb{Q}$, then its public price depends on the latent loading only through its public compression:*

$$C_t(a) = \mathbb{E}^{\mathbb{Q}}[X_a \mid \mathcal{H}_t] = C_{\mathcal{H}_t}^{\mathbb{Q}} L_{t,T(a)}^{\mathbb{Q}}(X_a).$$

Proof. The latent loading is $L_{t,T(a)}^{\mathbb{Q}}(X_a) = \mathbb{E}^{\mathbb{Q}}[X_a \mid \mathcal{K}_t]$. Since $\mathcal{H}_t \subseteq \mathcal{K}_t$, the tower property gives

$$\mathbb{E}^{\mathbb{Q}}[X_a \mid \mathcal{H}_t] = \mathbb{E}^{\mathbb{Q}}[\mathbb{E}^{\mathbb{Q}}[X_a \mid \mathcal{K}_t] \mid \mathcal{H}_t] = C_{\mathcal{H}_t}^{\mathbb{Q}} L_{t,T(a)}^{\mathbb{Q}}(X_a).$$

Thus the public price is the public compression of the latent loading. $\square$

The full-basket and multi-date results identify the information frontier of ideal complete claim families. They do not assert that listed-option datasets contain such manifolds. Empirical panels usually provide finite, noisy submanifolds. Later sections therefore keep finite-panel stability, source-family restrictions, and residual selection separate from full state recovery.

Definition 2.12 (Full terminal basket manifold). Fix $t < T$, and let $S_T = (S_T^1, \dots, S_T^N) \in \mathbb{R}_+^N$ be the terminal underlier vector. The full nonnegative basket-call manifold at (t, T) is

$$C_t(w, K; T) := \mathbb{E}^{\mathbb{Q}} \left[\frac{B_t}{B_T} (w^\top S_T - K)^+ \mid \mathcal{H}_t \right], \quad w \in \mathbb{R}_+^N, K \geq 0.$$

If signed spread baskets are admitted, w may range over $\mathbb{R}^N$ and K over $\mathbb{R}$.

Theorem 2.13 (Static basket-law recovery). Assume the conditional pricing law of S_T is supported on $\mathbb{R}_+^N$, and assume the public discount factor $P(t, T)$ is known at date t , equivalently the zero-coupon payoff at maturity T is included in the admitted public span. Write

$$D_{t,T} := \frac{B_t}{B_T}, \quad P(t, T) := \mathbb{E}^{\mathbb{Q}}[D_{t,T} \mid \mathcal{H}_t].$$

The full nonnegative basket-call manifold, augmented by this public discount coordinate, determines the T -forward conditional pricing law

$$\mu_{t,T}^T(\omega, A) := \frac{\mathbb{E}^{\mathbb{Q}}[D_{t,T} \mathbf{1}_{\{S_T \in A\}} \mid \mathcal{H}_t](\omega)}{P(t, T)(\omega)}$$

up to $\mathbb{Q}$ -null sets. Consequently it determines the date- t price of every bounded Borel European payoff $f(S_T)$ by integration against $\mu_{t,T}^T$. If $D_{t,T}$ is $\mathcal{H}_t$ -measurable, for example under deterministic discounting conditional on $\mathcal{H}_t$, then $\mu_{t,T}^T$ coincides with the raw conditional law $\text{Law}_{\mathbb{Q}}(S_T \mid \mathcal{H}_t)$.

Proof. The public discount coordinate supplies $P(t, T)$. After division by $P(t, T)$, each fixed w gives the full call surface of the T -forward pricing distribution of $w^\top S_T$. The Breeden–Litzenberger distributional derivative recovers the law of each nonnegative scalar projection under $\mu_{t,T}^T$ [18]. Consequently the manifold determines, pathwise in the conditioning variable, the T -forward conditional Laplace transform

$$L_{t,T}^T(\xi) = \int_{\mathbb{R}_+^N} e^{-\xi^\top x} \mu_{t,T}^T(dx), \quad \xi \in \mathbb{R}_+^N,$$

since $\xi = sw$ for some $s \geq 0$ and $w \in \mathbb{R}_+^N$. The multivariate Laplace transform on the positive orthant uniquely determines a probability law on $\mathbb{R}_+^N$. Prices of bounded European payoffs then follow from

$$\mathbb{E}^{\mathbb{Q}}[D_{t,T} f(S_T) \mid \mathcal{H}_t] = P(t, T) \int f(x) \mu_{t,T}^T(dx).$$

$\square$

Corollary 2.14 (Terminal-law maximality). Single-name vanilla surfaces recover only the one-dimensional T -forward pricing marginals of S_T . A finite family of basket surfaces recovers only the pricing laws of finitely many projections $w_a^\top S_T$. The full terminal basket manifold is maximal for terminal European payoffs at date T . Path-law information lies outside that terminal object. Knowing the terminal pricing law for each T does not by itself identify lead-lag structure, directed Volterra kernels, quadratic covariation paths, or the joint pricing law of $(S_{T_1}, \dots, S_{T_m})$ across future dates.

Proof. Theorem 2.13 identifies the full terminal pricing law from the full basket manifold. Restricting the manifold to single-name vanillas leaves only coordinate marginals; restricting it to finitely many baskets leaves only finitely many projection laws. None of these terminal objects specifies a coupling across future dates. Directed kernels, lead-lag effects, quadratic-covariation paths, and multi-date pricing laws are functions of that missing coupling, and therefore require the multi-date manifold of Theorem 2.16 or additional source restrictions. $\square$

Definition 2.15 (Multi-date linear option manifold). Fix $t < s_1 < \dots < s_m \leq T_*$. Let X_s be a vector of tradeable state coordinates, such as log-prices, forward returns, or centered variance factors. The multi-date linear option manifold is

$$C_t^{(s_1, \dots, s_m)}(\alpha, K) := \mathbb{E}^{\mathbb{Q}} \left[\frac{B_t}{B_{s_m}} \left(\sum_{k=1}^m \alpha_k^\top X_{s_k} - K \right)^+ \middle| \mathcal{H}_t \right],$$

for all admitted $\alpha = (\alpha_1, \dots, \alpha_m) \in \mathbb{R}^{mN}$ and strikes $K \in \mathbb{R}$.

Theorem 2.16 (Finite-dimensional path-law recovery). *For a fixed grid $s_1 < \dots < s_m$, assume the public discount factor $P(t, s_m)$ is known, or that the zero-coupon payoff at s_m is included in the admitted public span. The full multi-date linear option manifold determines the s_m -forward conditional pricing law of*

$$(X_{s_1}, \dots, X_{s_m}).$$

If these manifolds are observed for every finite grid in $[t, T_]$, and the recovered conditional pricing laws are projectively consistent and tight on the chosen path space, they determine the conditional path pricing law of $(X_s)_{s \in [t, T_*]}$. If all involved discount factors are $\mathcal{H}_t$ -measurable, the same statements hold for the raw $\mathbb{Q}$ -conditional laws.*

Proof. After normalization by the public discount factor $P(t, s_m)$, each α -slice gives a full call surface in K and therefore the s_m -forward pricing law of the scalar projection $\sum_k \alpha_k^\top X_{s_k}$. The Cramer–Wold theorem recovers the joint pricing law on $\mathbb{R}^{mN}$. Projective consistency across grids and tightness then give the conditional path pricing law by Kolmogorov extension, with the usual path-space regularity condition when continuous or càdlàg paths are imposed. Public discounting reduces this to raw conditional-law recovery. $\square$

Corollary 2.17 (Gaussian–Volterra covariance recovery). *Suppose X is conditionally Gaussian given $\mathcal{H}_t$ and admits the Volterra representation*

$$X_s = \int_t^s K(s, r) dW_r, \quad t \leq r \leq s \leq T_*,$$

with a square-integrable causal kernel $K(s, r)$. The multi-date manifold recovers the conditional covariance kernel

$$\Gamma(s, u) = \text{Cov}^{\mathbb{Q}}(X_s, X_u \mid \mathcal{H}_t) = \int_t^{s \wedge u} K(s, r) K(u, r)^\top dr.$$

The kernel K is identified up to orthogonal driver gauge. If a causal lower-triangular Volterra gauge with positive diagonal is imposed, the kernel is uniquely identified by the operator Cholesky factorization of Γ .

Proof. Theorem 2.16 recovers the conditional finite-dimensional laws of the process on every grid. For a conditionally Gaussian process, these laws are determined by their conditional means and covariance matrices; in the centered Volterra representation the mean is zero and the covariances are the displayed values of Γ . Two causal square-integrable kernels with the same covariance operator differ by an orthogonal transformation of the driving noise on the driver space. Imposing the causal lower-triangular gauge with positive diagonal selects the operator Cholesky factor and removes this driver gauge. $\square$

Theorem 2.18 (Numeraire invariance of local recovery). *Let N be a strictly positive traded numeraire and let $\mathbb{Q}^N$ be the associated equivalent measure. Suppose the money-market response*

map R_t and the N -numeraire response map R_t^N are related by a bounded invertible $\mathcal{H}_t$ -measurable linear transformation

$$R_t^N = T_t R_t$$

of the quote space. Then

$$D(R_t^N \circ \psi)(u_0) = T_t D(R_t \circ \psi)(u_0).$$

Local injectivity, quotient rank, finite-panel information rank, and overidentifying normal-space restrictions are invariant under the numeraire change. Inverse constants change only by the operator norms of T_t and T_t^{-1} .

Proof. The change-of-numeraire identity gives equality of arbitrage-free prices after the public linear quote transformation T_t . Since T_t is invertible, it neither creates nor removes null directions of the response derivative on the quotient tangent space. Left inverses transform by composition with T_t^{-1} , giving the norm statement. $\square$

3. THE RISK-NEUTRAL VOLTERRA–PERRON STATE

The state is introduced at the quotient level. Gaussian–Volterra assumptions enter later as one source family verifying the required response ranks.

Definition 3.1 (Risk-neutral Volterra–Perron state). The risk-neutral Volterra–Perron state at date t is

$$\Xi_t^{\mathbb{Q}} = (G_t^{\mathbb{Q}}, \mathfrak{H}_t^{\mathbb{Q}}, \rho_t^{\mathbb{Q}}, e_t^{\mathbb{Q}}, \ell_t^{\mathbb{Q}}, \varpi_t^{\mathbb{Q}}),$$

where:

- (1) $G_t^{\mathbb{Q}}$ is a directed Volterra operator on a Hilbert space $\mathcal{U}_t$;
- (2) $\mathfrak{H}_t^{\mathbb{Q}}$ is a rough maturity-scaling coordinate;
- (3) $(\rho_t^{\mathbb{Q}}, e_t^{\mathbb{Q}}, \ell_t^{\mathbb{Q}})$ is a right-left Perron spectral triple for a screened positive synchronization operator $A_t^{\mathbb{Q}}$, normalized by $\ell_t^{\mathbb{Q}}(e_t^{\mathbb{Q}}) = 1$;
- (4) $\varpi_t^{\mathbb{Q}} \in L^2(\mathbb{Q}, \mathcal{K}_t)$ is a centered synchronized latent stress coordinate;

A selected pricing state is

$$\Theta_t^{\mathbb{Q}} = (\Xi_t^{\mathbb{Q}}, \theta_t),$$

where θ_t is a residual-density selector used to price latent residual claims after the state has been recovered or imposed.

Definition 3.2 (Perron spectral projector). Let A be a screened positive synchronization operator with simple Perron eigenvalue ρ_A , right Perron vector e_A , and left Perron functional ℓ_A , normalized by $\ell_A(e_A) = 1$. The Perron spectral projector is

$$\Pi_A := e_A \otimes \ell_A, \quad \Pi_A x = e_A \ell_A(x).$$

The canonical Perron coordinate is (ρ_A, Π_A) . The triple (ρ_A, e_A, ℓ_A) is retained for computations, but the projector is the normalization-free coordinate of the Perron mode.

Proposition 3.3 (Gauge invariance of the Perron projector). The projector $\Pi_A = e_A \otimes \ell_A$ is invariant under reciprocal normalization changes

$$e_A \mapsto c e_A, \quad \ell_A \mapsto c^{-1} \ell_A, \quad c > 0.$$

If A is conjugated by an admissible hidden-basis change S , so that $\tilde{A} = SAS^{-1}$, then

$$\Pi_{\tilde{A}} = S \Pi_A S^{-1}.$$

Thus (ρ_A, Π_A) is the natural quotient coordinate of the Perron synchronization mode.

Proof. Reciprocal scaling gives $(c e_A) \otimes (c^{-1} \ell_A) = e_A \otimes \ell_A$. Under conjugacy, $S e_A$ and $\ell_A S^{-1}$ satisfy

$$S A S^{-1} (S e_A) = \rho_A S e_A, \quad (\ell_A S^{-1}) S A S^{-1} = \rho_A (\ell_A S^{-1}),$$

and $(\ell_A S^{-1})(S e_A) = 1$. The rank-one projector is therefore transported by $S \Pi_A S^{-1}$. $\square$

Definition 3.4 (Option pricing map). Let $\mathfrak{X}_t$ be an admissible state space for $\Theta_t^{\mathbb{Q}}$. The option pricing map is an adapted map

$$\mathfrak{P}_t : \mathfrak{X}_t \rightarrow \mathcal{Y}_t, \quad \Theta_t^{\mathbb{Q}} \mapsto \mathcal{M}_t^{\text{opt}}.$$

The map is public because $\mathfrak{P}_t(\Theta_t^{\mathbb{Q}})$ is $\mathcal{H}_t$ -measurable even when $\Theta_t^{\mathbb{Q}}$ contains $\mathcal{K}_t$ -measurable latent components.

Definition 3.5 (Conditional public response). When a source state contains latent $\mathcal{K}_t$ -measurable coordinates, the public response map is a conditional response. If $\tilde{\mathfrak{P}}_t(\Theta) \in L^1(\mathbb{Q}, \mathcal{K}_t; \mathcal{Y}_t)$ denotes the latent quote response generated before public compression, the public option manifold is

$$\mathfrak{P}_t(\Theta) := \mathbb{E}^{\mathbb{Q}}[\tilde{\mathfrak{P}}_t(\Theta) \mid \mathcal{H}_t].$$

If the recovered state coordinates are already represented by $\mathcal{H}_t$ -measurable canonical parameters, this convention reduces to the ordinary deterministic response chart. All public-indistinguishability and recovery statements below are statements about this $\mathcal{H}_t$ -measurable response, not about direct observation of hidden sample-path coordinates.

Remark 3.6. The recovery theorems identify the canonical class of $\Xi_t^{\mathbb{Q}}$. When the selector θ_t enters only through public-annihilating residual directions, admitted public option quotes leave it unidentified. A post-recovery rule must fix it, or an augmented manifold must add residual-price instruments.

Definition 3.7 (Calibration ledger). At date t , a finite option-manifold calibration separates four maps:

$$\begin{aligned} \text{public fit:} \quad & y_t \approx \mathfrak{P}_t(\Theta), \\ \text{source recovery:} \quad & y_t \mapsto [\Xi_t^{\mathbb{Q}}] \in \bar{\mathfrak{X}}_t, \\ \text{residual selection:} \quad & [\Xi_t^{\mathbb{Q}}] \mapsto \theta_t, \end{aligned}$$

and

$$\text{tradeable completion:} \quad X \mapsto \Pi_{\mathcal{S}_t} X.$$

The first is quote calibration in the public option panel. The second is source recovery on the canonical quotient. The third chooses a public-equivalent residual price inside the remaining envelope. The fourth is projection onto the finite public instrument span.

Remark 3.8 (Calibration, recovery, selection, and hedging). In a complete Black–Scholes-style finite model these four maps may collapse into one fitted parameter vector and one hedge. In the Volterra–Perron option manifold they are logically separate. A surface can be well calibrated without source recovery; a source can be recovered without identifying the residual selector; and a selected residual price can remain outside the tradeable span. The recovery theorems below concern the second map.

Definition 3.9 (Gauge equivalence). Two states $\Theta, \tilde{\Theta} \in \mathfrak{X}_t$ are gauge equivalent, written $\Theta \sim_g \tilde{\Theta}$, if they differ only by transformations that leave all latent loadings and all option-manifold prices unchanged, namely sign conventions for one-dimensional latent coordinates, bounded invertible basis changes in hidden factor coordinates, common rescalings absorbed by the Perron normalization $\ell(e) = 1$, and relabelings of unobserved driver bases. The canonical state space is the quotient

$$\bar{\mathfrak{X}}_t := \mathfrak{X}_t / \sim_g.$$

Remark 3.10. The quotient is essential. Perron eigenvectors, Volterra driver bases, and one-dimensional stress coordinates have unavoidable normalization choices. Recovery of the latent state can only mean recovery of the equivalence class $[\Xi_t^{\mathbb{Q}}] \in \bar{\mathfrak{X}}_t$.

Remark 3.11. The coordinate $\varpi_t^{\mathbb{Q}}$ is latent in parametrization but recoverable only through its loading, conditional moment, or canonical horizon lift encoded by the option manifold. The recovery statements do not assert observation of an arbitrary future sample-path realization. When a terminal stress payoff is needed, $\varpi_{t,T}^{\mathbb{Q}}$ denotes the horizon-lifted synchronized stress direction generated by the recovered source chart.

Assumption 3.12 (Local gauge slice). Near a reference state Ξ_0 , the gauge transformations form a C^1 local group $\mathcal{G}_t$ acting on the source chart, and the public response map is constant on $\mathcal{G}_t$ -orbits. The orbit tangent

$$E_g := T_{\Xi_0}(\mathcal{G}_t \cdot \Xi_0)$$

is closed and admits a closed complement E_c , so $E = E_g \oplus E_c$. A gauge-normalized chart is a local slice $U_c \subset E_c$ meeting each nearby orbit once. The quotient tangent space is identified with E_c .

Remark 3.13 (Local gauge fixing). Assumption 3.12 is intentionally local. It gives a cross-section near the recovered state used in the inverse-function and stability arguments. It does not assert the existence of a global, singularity-free section of the full orbit space $\mathfrak{X}_t/\mathcal{G}_t$. Global stabilizers, orientation changes, eigenvector normalization seams, or nontrivial orbit topology may force an atlas of slices rather than a single canonical representative. All recovery statements below are therefore local or patched through the recovery atlas of Section 9; the quotient state is global only up to the declared gauge equivalence.

Definition 3.14 (Gauge-normalized source chart). A source chart is gauge-normalized near Ξ_0 if each state is represented by canonical coordinates

$$u = (u_{\mathfrak{H}}, u_G, u_\rho, u_e, u_P) \in E_{\mathfrak{H}} \oplus E_G \oplus E_\rho \oplus E_e \oplus E_P,$$

corresponding to

$$(\mathfrak{H}^{\mathbb{Q}}, G^{\mathbb{Q}}, \rho^{\mathbb{Q}}, e^{\mathbb{Q}}, \ell^{\mathbb{Q}}, P^{\mathbb{Q}}),$$

with the following conventions:

- (1) $\ell^{\mathbb{Q}}(e^{\mathbb{Q}}) = 1$, $e^{\mathbb{Q}}$ lies in a fixed positive cone, and the remaining Perron normalization is absorbed into $P^{\mathbb{Q}}$;
- (2) the latent stress coordinate $P^{\mathbb{Q}}$ is centered and has a fixed sign convention against a nonzero reference functional;
- (3) hidden Volterra driver bases are represented in a fixed orthonormal or triangular gauge.

The block u_e should be read as the normalized Perron-projector block: the left functional $\ell^{\mathbb{Q}}$ is induced by the screened operator and the normalization $\ell^{\mathbb{Q}}(e^{\mathbb{Q}}) = 1$ inside the chosen slice. A redundant chart may include an explicit u_ℓ coordinate, but the quotient tangent then removes the normalization direction and returns the same canonical block. The corresponding quotient tangent space is the direct sum of these gauge-normalized coordinate blocks.

4. A VOLTERRA–PERRON SOURCE CHART

The recovery theorems are stated for general source charts. This section fixes a finite Volterra–Perron chart in which the response ranks can be read block by block. Rough maturity scaling carries Volterra source geometry; basket and carrier responses carry covariance routes; the screened route matrix carries the Perron synchronization coordinate. Residual selectors are added only afterward as public-equivalent pricing extensions, unless public completion or residual instruments make them observable.

Definition 4.1 (Volterra route kernels). Fix N underliers and a maximum maturity $\bar{\tau} > 0$. Conditional on $\mathcal{H}_t$, let

$$B_{t+s}^{\mathbb{Q}} = (B_{1,t+s}^{\mathbb{Q}}, \dots, B_{N,t+s}^{\mathbb{Q}})$$

be a risk-neutral Gaussian driver vector with conditionally independent standard components. A nontrivial instantaneous driver covariance can be included as an additional gauge-normalized matrix coordinate; the displayed route formula fixes the independent-driver gauge. The centered latent forward-variance stress of name i over future time $s \in [0, \bar{\tau}]$ is represented by

$$Y_{i,t}^{\mathbb{Q}}(s) = \sum_{j=1}^N \int_0^s K_{ij,t}^{\mathbb{Q}}(s-r) dB_{j,t+r}^{\mathbb{Q}}.$$

The route from source j to target i has kernel

$$K_{ij,t}^{\mathbb{Q}}(x) = a_{ij,t} x^{H_{ij,t}^{\mathbb{Q}} - \frac{1}{2}} \exp(-\beta_{ij,t} x) \left(1 + \sum_{k=1}^{K_0} c_{ij,k,t} x^k \right), \quad 0 < x \leq \bar{\tau},$$

where $H_{ij,t}^{\mathbb{Q}} \in (0, 1)$, $\beta_{ij,t} \geq 0$, and the polynomial modulation is bounded away from zero on the recovery neighborhood. The directed Volterra operator $G_t^{\mathbb{Q}}$ is the matrix of off-diagonal route coordinates

$$G_{ij,t}^{\mathbb{Q}} = (a_{ij,t}, H_{ij,t}^{\mathbb{Q}}, \beta_{ij,t}, c_{ij,1,t}, \dots, c_{ij,K_0,t}), \quad i \neq j.$$

The observable covariance routes generated by these kernels are explicit. For $\tau \leq \bar{\tau}$, define

$$V_{ij,t}^{\mathbb{Q}}(\tau) := \text{Cov}^{\mathbb{Q}} \left(\int_0^\tau Y_{i,t}^{\mathbb{Q}}(s) ds, \int_0^\tau Y_{j,t}^{\mathbb{Q}}(s) ds \middle| \mathcal{H}_t \right).$$

Equivalently,

$$V_{ij,t}^{\mathbb{Q}}(\tau) = \sum_{q=1}^N \int_0^\tau \left(\int_r^\tau K_{iq,t}^{\mathbb{Q}}(s-r) ds \right) \left(\int_r^\tau K_{jq,t}^{\mathbb{Q}}(s-r) ds \right) dr.$$

Lemma 4.2 (Analytic leading power of a route covariance). *Suppose the q -source routes entering targets i and j satisfy, as $x \downarrow 0$,*

$$K_{iq,t}^{\mathbb{Q}}(x) = a_{iq,t} x^{H_{iq,t}^{\mathbb{Q}} - \frac{1}{2}} (1 + O(x)), \quad K_{jq,t}^{\mathbb{Q}}(x) = a_{jq,t} x^{H_{jq,t}^{\mathbb{Q}} - \frac{1}{2}} (1 + O(x)),$$

with nonzero amplitudes. Then the corresponding route contribution to $V_{ij,t}^{\mathbb{Q}}(\tau)$ has the short-maturity expansion

$$V_{ij,t}^{(q),\mathbb{Q}}(\tau) = A_{ijq,t} \tau^{H_{iq,t}^{\mathbb{Q}} + H_{jq,t}^{\mathbb{Q}} + 2} (1 + O(\tau)),$$

where

$$A_{ijq,t} = \frac{a_{iq,t} a_{jq,t}}{(H_{iq,t}^{\mathbb{Q}} + \frac{1}{2})(H_{jq,t}^{\mathbb{Q}} + \frac{1}{2})(H_{iq,t}^{\mathbb{Q}} + H_{jq,t}^{\mathbb{Q}} + 2)}$$

up to the displayed normalization of the leading kernel amplitudes.

Proof. For $0 \leq r \leq \tau$, integrate the leading kernel expansion in the variable $s - r$:

$$\int_r^\tau K_{iq,t}^{\mathbb{Q}}(s-r) ds = \frac{a_{iq,t}}{H_{iq,t}^{\mathbb{Q}} + \frac{1}{2}} (\tau - r)^{H_{iq,t}^{\mathbb{Q}} + \frac{1}{2}} (1 + O(\tau - r)).$$

The same expression holds for the j -route. Their product has leading power $(\tau - r)^{H_{iq,t}^{\mathbb{Q}} + H_{jq,t}^{\mathbb{Q}} + 1}$. Hence

$$\int_0^\tau (\tau - r)^{H_{iq,t}^{\mathbb{Q}} + H_{jq,t}^{\mathbb{Q}} + 1} dr = \frac{\tau^{H_{iq,t}^{\mathbb{Q}} + H_{jq,t}^{\mathbb{Q}} + 2}}{H_{iq,t}^{\mathbb{Q}} + H_{jq,t}^{\mathbb{Q}} + 2},$$

and the displayed coefficient follows after multiplying by the two one-step integration constants. $\square$

Roughness and directed-route coefficients therefore enter the public manifold through analytic maturity powers and logarithmic derivatives. The finite dictionary gives a sufficient response family for exact finite-dimensional recovery. The lemma supplies the local invariant used when the full covariance route is not assumed to lie exactly in a finite power-exponential span.

Definition 4.3 (Rough power-exponential dictionary). On a maturity interval $(0, \bar{\tau})$, a finite Volterra response dictionary consists of functions of the form

$$\tau^\alpha e^{-\beta\tau} p(\tau), \quad \alpha > -1, \quad \beta \geq 0,$$

where p is a polynomial of degree at most K . Dictionary atoms are grouped by their pairs (α, β) . A dictionary used for identification is taken in reduced form. After expanding each polynomial, the scalar atoms

$$\{\tau^{\alpha_\ell + k} e^{-\beta_\ell \tau} : 0 \leq k \leq \deg p_\ell\}$$

are pairwise distinct and linearly independent on every open subinterval. A convenient sufficient check is that, within each fixed β , no two exponents created by the admitted polynomial shifts coincide. This removes notational duplicates such as $\tau^\alpha \tau$ and $\tau^{\alpha+1}$.

Theorem 4.4 (Power-exponential route identifiability). *Let*

$$f(\tau) = \sum_{\ell=1}^m \tau^{\alpha_\ell} e^{-\beta_\ell \tau} p_\ell(\tau), \quad 0 < \tau < \bar{\tau},$$

where the pairs $(\alpha_\ell, \beta_\ell)$ are distinct, each p_ℓ is a polynomial of known maximum degree, and the expanded scalar atoms form a reduced dictionary in the sense of Definition 4.3. If f vanishes on a nonempty open interval, then every expanded scalar coefficient is zero; equivalently every reduced dictionary coefficient vanishes. Hence the coefficient vector of a finite Volterra response dictionary is identified by the continuum maturity response. There is also a nonzero analytic finite-panel determinant whose nonvanishing identifies the same coefficient vector, and the deficient panels form a zero set with empty interior.

Proof. Expand each polynomial and collect the response in scalar atoms $\tau^{\gamma_r} e^{-\beta_r \tau}$. These atoms are real analytic on $(0, \infty)$. If the response vanishes on a nonempty open subinterval, the identity theorem gives vanishing on the whole positive half-line. Among the remaining nonzero coefficients, choose the smallest exponential rate β_* . Multiplying by $e^{\beta_* \tau}$ and letting $\tau \rightarrow \infty$ removes all larger-rate groups exponentially. Inside the β_* -group, distinct powers are separated by dividing successively by the largest remaining power of τ . Its coefficient must be zero; iteration eliminates the group. Repeating over the finite list of exponential rates eliminates all expanded scalar coefficients. Thus the continuum maturity response is injective on the reduced coefficient space.

It remains to pass from the continuum response to a finite panel. Let E be the finite coefficient space and let $\varepsilon_\tau : E \rightarrow \mathbb{R}$ be evaluation at maturity τ . If no finite family of evaluations separated E , the span of $\{\varepsilon_\tau : \tau \in I\}$ in E^* would be a proper subspace. A nonzero vector in its annihilator would generate a nonzero dictionary combination vanishing on all of I , contradicting the continuum injectivity. Thus some finite panel is injective. A full-rank square minor of the associated collocation matrix is nonzero at that panel. Since the entries are analytic in the maturity coordinates, this minor is a nonzero analytic determinant, and deficient panels lie in its zero set, which has empty interior. $\square$

Option-manifold recovery of a route coordinate is not high-frequency market admissibility. The power-exponential chart may recover smooth or semi-smooth option-route coordinates on the public surface. Rough market admissibility is the smaller source-screened class with local high-frequency normalization and $H_{ij} < 1/2$. Recovered option routes enter the rough market Perron layer only after that additional admissibility condition is imposed.

Corollary 4.5 (Rough exponent as a short-maturity invariant). *If a route response has leading form*

$$f(\tau) = A \tau^\alpha e^{-\beta \tau} (1 + c_1 \tau + O(\tau^2)), \quad A \neq 0,$$

then

$$\alpha = \lim_{\tau \downarrow 0} \tau \frac{f'(\tau)}{f(\tau)}$$

and

$$\beta - c_1 = - \lim_{\tau \downarrow 0} \left(\frac{f'(\tau)}{f(\tau)} - \frac{\alpha}{\tau} \right).$$

Thus the rough exponent is an observable short-maturity invariant whenever the leading route is separated.

Proof. Since $A \neq 0$, the response has fixed sign for all sufficiently small positive maturities. On that interval,

$$\log |f(\tau)| = \log |A| + \alpha \log \tau - \beta \tau + c_1 \tau + O(\tau^2).$$

Differentiation gives

$$\frac{f'(\tau)}{f(\tau)} = \frac{\alpha}{\tau} - \beta + c_1 + O(\tau),$$

and the two limits follow. $\square$

Definition 4.6 (Surface coordinates generated by covariance routes). Fix a pricing chart $\chi \in \{\text{B}, \text{N}\}$. Let $F_{i,t,T}$ be the public forward of name i to $T = t + \tau$, and let an index or basket u have holdings $w^u = (w_1^u, \dots, w_N^u)$ and forward $F_{u,t,T} = \sum_i w_i^u F_{i,t,T}$. The route covariance matrix $V_t^\mathbb{Q}(\tau)$ is written in normalized forward-return coordinates. The chart-specific exposure vector is

$$\alpha_i^{u,\text{B}}(t, T) := \frac{w_i^u F_{i,t,T}}{F_{u,t,T}}, \quad \alpha_i^{u,\text{N}}(t, T) := w_i^u F_{i,t,T}.$$

For a single name these formulas are read with the corresponding unit holding. The local chart-specific total-variance coordinate is

$$\text{TV}_{u,t}^\chi(\tau) = \text{TV}_{u,t}^{\chi,0}(\tau) + \sum_{i,j=1}^N \alpha_i^{u,\chi}(t, T) \alpha_j^{u,\chi}(t, T) V_{ij,t}^\mathbb{Q}(\tau).$$

In the Black chart TV^B is dimensionless lognormal-forward total variance, while in the Bachelier chart TV^N is normal-forward total variance in squared cash-forward units. The corresponding level coordinates are

$$\text{IVL}_{u,t}^\text{B}(\tau) = \sqrt{\text{TV}_{u,t}^\text{B}(\tau)/\tau}, \quad \text{NVL}_{u,t}^\text{N}(\tau) = \sqrt{\text{TV}_{u,t}^\text{N}(\tau)/\tau}.$$

The local moneyness expansion of the surface in chart χ is

$$v_{u,t}^\chi(\tau, m^\chi) = \text{TV}_{u,t}^\chi(\tau) + \text{SKW}_{u,t}^\chi(\tau) m^\chi + \frac{1}{2} \text{CURV}_{u,t}^\chi(\tau) (m^\chi)^2 + O((m^\chi)^3),$$

where the downside-skew and curvature coordinates are finite linear functionals of $\{V_{ij,t}^\mathbb{Q}(\tau)\}_{i,j}$, for example

$$\text{SKW}_{u,t}^\chi(\tau) = s_{u,t}^{\chi,0}(\tau) + \sum_{i,j} b_{u,ij,t}^{\chi,\text{skw}}(\tau) V_{ij,t}^\mathbb{Q}(\tau).$$

For a call or put coordinate $a = (u, \tau, m^\chi, \kappa, \chi)$, the local pricing map is obtained by substituting $v_{u,t}^\chi(\tau, m^\chi)$ in the corresponding no-arbitrage one-period price chart:

$$C_t^\Theta(a) = \text{B}_\kappa^\chi(D_{t,T}, F_{u,t,T}, K^\chi(m^\chi), v_{u,t}^\chi(\tau, m^\chi)).$$

The construction is a local response chart around an arbitrage-free public surface. It does not assert a global parametrization of every strike and maturity. The recovery neighborhood is chosen inside the admitted no-arbitrage chart, so positivity, monotonicity, convexity, and calendar constraints remain valid on the finite or continuum panel under consideration. The two total-variance coordinates are chart-specific objects. A lognormal total variance and a normal-forward variance are not numerically interchangeable without an additional conversion model. The rank statements below use only the local price chart and its nonzero variance vega.

Definition 4.7 (Black and Bachelier price charts). Let $D > 0$ be a discount factor, $F > 0$ a forward, K a strike, and $v > 0$ a total-variance coordinate in the chosen chart. These are the Bachelier normal-forward and Black–Scholes–Merton lognormal-forward charts [19, 20, 21]. Let Φ_0 and φ_0 denote the standard normal distribution function and density. In the Black chart,

$$d_\pm = \frac{\log(F/K)}{\sqrt{v}} \pm \frac{1}{2} \sqrt{v},$$

$$\text{B}_C^\text{B}(D, F, K, v) = D(F\Phi_0(d_+) - K\Phi_0(d_-)), \quad \text{B}_P^\text{B}(D, F, K, v) = D(K\Phi_0(-d_-) - F\Phi_0(-d_+)).$$

In the Bachelier chart,

$$d = \frac{F - K}{\sqrt{v}},$$

$$\begin{aligned}\mathcal{B}_C^N(D, F, K, v) &= D((F - K)\Phi_0(d) + \sqrt{v}\varphi_0(d)), \\ \mathcal{B}_P^N(D, F, K, v) &= D((K - F)\Phi_0(-d) + \sqrt{v}\varphi_0(d)).\end{aligned}$$

Both charts satisfy put–call parity

$$\mathcal{B}_C^X(D, F, K, v) - \mathcal{B}_P^X(D, F, K, v) = D(F - K).$$

Lemma 4.8 (Variance vegas). *In the Black chart,*

$$\partial_v \mathcal{B}_C^B = \partial_v \mathcal{B}_P^B = \frac{DF\varphi_0(d_+)}{2\sqrt{v}}.$$

In the Bachelier chart,

$$\partial_v \mathcal{B}_C^N = \partial_v \mathcal{B}_P^N = \frac{D\varphi_0(d)}{2\sqrt{v}}.$$

Proof. In the Black chart,

$$\partial_v d_+ = -\frac{\log(F/K)}{2v^{3/2}} + \frac{1}{4\sqrt{v}}, \quad \partial_v d_- = \partial_v d_+ - \frac{1}{2\sqrt{v}}.$$

The density identity

$$F\varphi_0(d_+) = K\varphi_0(d_-)$$

cancels the $d_{\pm}$ -terms generated by differentiating the two distribution functions. The remaining Black contribution is

$$\frac{DF\varphi_0(d_+)}{2\sqrt{v}}.$$

In the Bachelier chart, $d = (F - K)/\sqrt{v}$. The derivatives of $(F - K)\Phi_0(d)$ and $\sqrt{v}\varphi_0(d)$ cancel except for $\varphi_0(d)/(2\sqrt{v})$. Put prices have the same variance derivative by put–call parity, since $D(F - K)$ does not depend on v . $\square$

Theorem 4.9 (Pricing-chart rank invariance). *Let $u \mapsto v(u) \in \mathbb{R}^M$ be a finite vector of chart-specific total-variance coordinates, and suppose public discount factors, forwards, strikes, maturities, and option types are fixed. Let*

$$C_r(u) = \mathcal{B}_{\kappa_r}^{X_r}(D_r, F_r, K_r, v_r(u)), \quad r = 1, \dots, M.$$

If the variance vegas $w_r := \partial_v \mathcal{B}_{\kappa_r}^{X_r}(D_r, F_r, K_r, v_r(u_0))$ satisfy $0 < \underline{w} \leq |w_r| \leq \bar{w} < \infty$ at u_0 , then

$$DC(u_0) = W Dv(u_0), \quad W = \text{diag}(w_1, \dots, w_M).$$

Consequently $DC(u_0)$ and $Dv(u_0)$ have the same rank, the same kernel, and comparable nonzero singular values with constants depending only on $\underline{w}$ and $\bar{w}$.

Proof. The chain rule and Lemma 4.8 give the diagonal factorization. Multiplying by an invertible diagonal matrix preserves rank and kernel. The singular-value comparison follows from

$$\underline{w} \|Dv(u_0)h\| \leq \|DC(u_0)h\| \leq \bar{w} \|Dv(u_0)h\|.$$

$\square$

Lemma 4.10 (Parity projection and no-arbitrage residuals). *For fixed public D, F, K , the call–put parity row*

$$C - P - D(F - K) = 0$$

is a static public no-arbitrage restriction, not an independent Volterra–Perron source response. Recovery ranks may therefore be computed after projecting quoted call–put pairs onto the parity-consistent subspace. The orthogonal distance to that subspace, in the quote norm of Definition 2.2, is a quote-quality or no-arbitrage-cleaning statistic rather than evidence for a latent residual kernel.

Proof. The identity is deterministic once D, F, K are fixed. It does not depend on the latent source chart, hence its differential annihilates all source directions. Projecting away the parity row removes a redundant public coordinate and leaves the source-response rank unchanged. A parity violation is therefore a violation of the admitted public price cone, not a new latent pricing direction. $\square$

Definition 4.11 (Screened synchronization and carrier response). Fix a synchronization maturity τ_* . Let $B_t^\mathbb{Q}$ denote the signed directed Volterra route matrix in the chosen source chart. Its entries are the active directed route coefficients whose public covariance-route images are $V_{ij,t}^\mathbb{Q}(\tau)$. The screened risk-neutral synchronization matrix is the positive option-manifold image

$$A_t^\mathbb{Q} = \Phi_t(B_t^\mathbb{Q}).$$

On the active screened stratum, Φ_t is the differentiable map with off-diagonal entries

$$A_{ij,t}^\mathbb{Q} = q_{ij,t} \frac{(V_{ij,t}^\mathbb{Q}(\tau_*))_+}{\sqrt{(V_{ii,t}^\mathbb{Q}(\tau_*) + \varepsilon)(V_{jj,t}^\mathbb{Q}(\tau_*) + \varepsilon)}}, \quad i \neq j,$$

and $A_{ii,t}^\mathbb{Q} = 0$. Here $q_{ij,t} \geq 0$ is a source-screening weight and $\varepsilon > 0$ is fixed. The normalization by the diagonal variance terms makes $A_t^\mathbb{Q}$ a dimensionless screened one-step synchronization operator; Perron critical surfaces therefore refer to this normalized object rather than to an arbitrary scale of option variance. Passing from $B_t^\mathbb{Q}$ to $A_t^\mathbb{Q}$ loses sign information and retains the positive synchronization geometry observed on the option manifold. Here the screened synchronization operator is the finite matrix $A_t^\mathbb{Q} \in \mathbb{R}_+^{N \times N}$ on the active underlier set. Infinite-dimensional Krein–Rutman variants would require compact strong positivity assumptions and are not used here. If $A_t^\mathbb{Q}$ is irreducible with simple Perron root, then

$$A_t^\mathbb{Q} e_t^\mathbb{Q} = \rho_t^\mathbb{Q} e_t^\mathbb{Q}, \quad \ell_t^\mathbb{Q} A_t^\mathbb{Q} = \rho_t^\mathbb{Q} \ell_t^\mathbb{Q}, \quad \ell_t^\mathbb{Q}(e_t^\mathbb{Q}) = 1.$$

Given carrier portfolios $b^1, \dots, b^r \in \mathbb{R}^N$, their synchronization response vectors are

$$C_k(\tau) = \sum_{i,j} b_i^k b_j^k V_{ij,t}^\mathbb{Q}(\tau), \quad k = 1, \dots, r,$$

and the carrier Gram matrix on a maturity grid $\mathcal{T}$ is

$$\Gamma_{\text{car},kl} = \sum_{\tau \in \mathcal{T}} C_k(\tau) C_l(\tau).$$

Carrier absorption is nondegenerate when $\Gamma_{\text{car}} \succeq \kappa_{\text{car}} I$.

Lemma 4.12 (Screened-edge derivative). *On an active screened stratum where $V_{ij,t}^\mathbb{Q}(\tau_*) > 0$ and $q_{ij,t}$ is differentiable, write*

$$A_{ij} = q_{ij} \frac{V_{ij}}{D_{ij}}, \quad D_{ij} := \sqrt{(V_{ii} + \varepsilon)(V_{jj} + \varepsilon)}.$$

Then

$$\delta A_{ij} = \frac{q_{ij}}{D_{ij}} \delta V_{ij} - \frac{A_{ij}}{2} \left(\frac{\delta V_{ii}}{V_{ii} + \varepsilon} + \frac{\delta V_{jj}}{V_{jj} + \varepsilon} \right) + \frac{V_{ij}}{D_{ij}} \delta q_{ij}.$$

If q_{ij} is fixed, the last term is absent. The fixed-sign margin in Proposition 4.16 is exactly the condition that keeps this derivative away from the kink of the positive-part screening map.

Proof. On the active stratum the positive-part map has derivative one, so $A_{ij} = q_{ij} V_{ij} D_{ij}^{-1}$. Differentiating this quotient gives

$$\delta A_{ij} = \frac{q_{ij}}{D_{ij}} \delta V_{ij} + \frac{V_{ij}}{D_{ij}} \delta q_{ij} - A_{ij} \frac{\delta D_{ij}}{D_{ij}}.$$

Since

$$\frac{\delta D_{ij}}{D_{ij}} = \frac{1}{2} \left(\frac{\delta V_{ii}}{V_{ii} + \varepsilon} + \frac{\delta V_{jj}}{V_{jj} + \varepsilon} \right),$$

the displayed formula follows. $\square$

Definition 4.13 (Interval screened matrix). Suppose finite-panel recovery or bid–ask uncertainty gives elementwise bounds on the nonnegative screened synchronization matrix,

$$L_{ij} \leq A_{ij} \leq U_{ij}, \quad L, U \geq 0.$$

The compatible matrix set is

$$\mathcal{A}[L, U] := \{A \geq 0 : L \leq A \leq U \text{ elementwise}\}.$$

Perron–Frobenius terminology and matrix perturbation conventions are used in their finite-dimensional sense [22, 23].

Theorem 4.14 (Certified Perron phase bounds). *For every $A \in \mathcal{A}[L, U]$,*

$$\rho(L) \leq \rho(A) \leq \rho(U),$$

where $\rho(\cdot)$ denotes spectral radius. For any strictly positive vector x ,

$$\min_i \frac{(Lx)_i}{x_i} \leq \rho(A) \leq \max_i \frac{(Ux)_i}{x_i}.$$

Consequently, if $\rho(U) < \rho_{\text{crit}}$, the recovered state is certified below the synchronization threshold for every admissible screened matrix; if $\rho(L) > \rho_{\text{crit}}$, it is certified above the threshold. Otherwise the phase label is not identified at the available quote precision.

Proof. Elementwise order gives $L \leq A \leq U$. Monotonicity of the Perron root for nonnegative matrices therefore yields $\rho(L) \leq \rho(A) \leq \rho(U)$. For any strictly positive x , the Collatz–Wielandt inequalities give

$$\min_i \frac{(Lx)_i}{x_i} \leq \rho(L), \quad \rho(U) \leq \max_i \frac{(Ux)_i}{x_i},$$

and hence the displayed bound for every compatible A . The phase certification is the direct comparison of these bounds with ρ_{crit} . $\square$

Definition 4.15 (Bounded residual selector). Let $R^1, \dots, R^m$ be bounded residual directions satisfying

$$\mathbb{E}^{\mathbb{Q}}[R^r \mid \mathcal{H}_t] = 0, \quad \mathbb{E}^{\mathbb{Q}}[R^r X_a \mid \mathcal{H}_t] = 0 \quad \text{for every admitted public payoff } X_a.$$

For recovered state coordinates

$$z_t = \left(\varpi_t^{\mathbb{Q}}, \log \rho_t^{\mathbb{Q}}, \bar{\mathfrak{H}}_t^{\mathbb{Q}}, \sum_{i,j} \ell_{i,t}^{\mathbb{Q}} G_{ij,t}^{\mathbb{Q}} e_{j,t}^{\mathbb{Q}} \right),$$

define

$$\theta_{r,t} = \bar{\theta}_{r,t} \tanh\left(\eta_{r,0,t} + \eta_{r,t}^{\top} z_t\right), \quad \sum_{r=1}^m |\bar{\theta}_{r,t}| \|R^r\|_{\infty} < 1.$$

The associated public-equivalent density family is

$$Z_{t,S}^{\theta} = 1 + \sum_{r=1}^m \theta_{r,t} R^r.$$

This form keeps positivity and public equivalence explicit while allowing residual prices to depend on recovered Perron stress, spectral radius, roughness, and directed-route pressure. The rule operates after state recovery; it becomes a public-response coordinate only when augmented residual instruments are admitted.

Proposition 4.16 (Functional-form rank criterion). *In the source chart of Definitions 4.1–4.15, suppose:*

- (1) *the maturity-response atoms used for recovery form a reduced, linearly independent power-exponential dictionary after the chosen gauge normalizations;*
- (2) *the corresponding route amplitudes $A_{ijq,t}$ are nonzero;*
- (3) *the active screened edges satisfy a fixed-sign margin*

$$|V_{ij,t}^{\mathbb{Q}}(\tau_*)| \geq v_* > 0$$

on the recovery neighborhood, while inactive edges have fixed screening weight $q_{ij,t} = 0$;

- (4) *the index and carrier weight systems separate the cross-covariance routes entering $\{V_{ij,t}^{\mathbb{Q}}\}$;*

(5) the screened synchronization matrix has a simple Perron root and the carrier Gram matrix has a positive lower bound.

Then the public observable response derivative for the canonical state is block triangular with full-rank diagonal blocks. Consequently the block-rank observability criterion and the conditional recovery theorem apply. If an augmented manifold also contains residual-price instruments, the same criterion applies to the additional selector block; for a purely public option manifold the selector block is zero.

Proof. The route kernels generate analytic maturity functions in a reduced power-exponential dictionary. These atoms are linearly independent on any open maturity interval, and on a sufficiently resolving finite panel by Theorem 8.16. The nonzero route amplitudes make the directed Volterra block visible. The fixed-sign margin keeps the positive-part screening map differentiable on the active stratum and fixes inactive edges outside the quotient tangent space. The index and carrier weight conditions make the linear map from cross-covariance routes to index/carrier responses injective on the admitted source coordinates. The simple, well-conditioned Perron eigenvalue gives differentiable recovery of the spectral triple, while the carrier Gram lower bound gives a bounded inverse for carrier absorption coordinates.

After gauge directions have been removed, any public kernel must lie in one of the diagonal blocks or in a selector coordinate not captured by the pure public response. The stated rank conditions exclude the diagonal-block kernels. Public-annihilating residual selectors remain outside the pure public response derivative unless residual instruments are appended. The block operator is therefore triangular with full-rank diagonal blocks, and Theorem 8.13 applies. $\square$

Theorem 4.17 (Standard-model recovery mechanism). *In the source chart of this section, suppose the rank conditions of Proposition 4.16 hold on a recovery neighborhood. Then the admitted option manifold locally recovers the canonical source state*

$$(\mathfrak{H}^{\mathbb{Q}}, G^{\mathbb{Q}}, \rho^{\mathbb{Q}}, e^{\mathbb{Q}}, \ell^{\mathbb{Q}}, P^{\mathbb{Q}})$$

up to gauge. The inverse is assembled in four conceptual blocks:

maturity scaling $\implies \mathfrak{H}^{\mathbb{Q}}$ and the Volterra source dictionary,

basket/carrier polarization $\implies G^{\mathbb{Q}}$ and the active covariance routes,

screened synchronization $\implies (\rho^{\mathbb{Q}}, e^{\mathbb{Q}}, \ell^{\mathbb{Q}}, P^{\mathbb{Q}})$,

residual completion or selector choice $\implies \theta$ only outside the pure public response.

Consequently the pure public option manifold can recover the Volterra–Perron source state on this stratum. Residual-selector recovery requires public completion, augmented residual instruments, or a separate canonical selector rule.

Proof. The maturity and route blocks are injective by the analytic dictionary and carrier-separation conditions in Proposition 4.16. On the simple-Perron stratum, the screened synchronization block has a differentiable spectral chart whose derivative is controlled by the reduced resolvent; the carrier Gram lower bound gives a bounded inverse for the carrier absorption coordinates. Hence the quotient derivative has the block-triangular full-rank form of Definition 8.12. Theorem 8.13 gives a bounded left inverse, and the Volterra–Perron option-manifold equivalence theorem gives local recovery of the canonical source state. Because the residual selector lies outside the pure public response block, it is recovered only under the additional residual-completion alternatives stated in the theorem. $\square$

5. RISK-NEUTRAL MARKET PHASES AND TRANSITION SURFACES

The market-phase theory of [24] is imposed here on the recovered risk-neutral quotient. Phases are not exogenous regimes. They are functions of the recovered option-manifold state, recording whether public prices represent local risk, synchronized carrier-absorbed risk, residual stress, or synchronized risk without an adequate carrier span.

Definition 5.1 (Risk-neutral phase coordinates). In the source chart of Section 4, use the Perron spectral radius $\rho_t^{\mathbb{Q}}$ and define the remaining phase coordinates

$$K_t^{\text{car}} := \lambda_{\min}(\Gamma_{\text{car},t}),$$

$$R_t^{\theta} := \left(\mathbb{E}^{\mathbb{Q}} \left[\left(\sum_{r=1}^m \theta_{r,t} R^r \right)^2 \middle| \mathcal{H}_t \right] \right)^{1/2} = \left(\theta_t^{\top} G_t^R \theta_t \right)^{1/2}, \quad (G_t^R)_{rs} := \mathbb{E}^{\mathbb{Q}}[R^r R^s \mid \mathcal{H}_t],$$

and

$$D_t^{\mathbb{Q}} := \sum_{i,j} \ell_{i,t}^{\mathbb{Q}} q_{ij,t} a_{ij,t} e_{j,t}^{\mathbb{Q}}.$$

Here $\rho_t^{\mathbb{Q}}$ measures synchronized spectral amplification, K_t^{car} measures carrier absorption, R_t^{θ} measures residual-kernel pressure with the full conditional residual Gram matrix, and $D_t^{\mathbb{Q}}$ measures directed Perron-weighted route pressure.

Remark 5.2. For smooth phase weights, the recovery neighborhood is chosen so that the smallest carrier Gram eigenvalue is simple and $R_t^{\theta} > 0$, or else the smooth scores below are read with the regularized substitutes $\log \det(\Gamma_{\text{car},t} + \varepsilon I)$ and $\theta_t^{\top} G_t^R \theta_t$. The hard phase labels only require the continuous nonnegative coordinates.

Definition 5.3 (Hard phase regions). Fix public thresholds

$$\rho_* > 0, \quad K_* > 0, \quad R_* > 0.$$

The recovered state is in the local phase if

$$\rho_t^{\mathbb{Q}} < \rho_*, \quad R_t^{\theta} < R_*.$$

It is in the synchronized-carrier phase if

$$\rho_t^{\mathbb{Q}} \geq \rho_*, \quad K_t^{\text{car}} \geq K_*, \quad R_t^{\theta} < R_*.$$

It is in the residual-stress phase if

$$\rho_t^{\mathbb{Q}} \geq \rho_*, \quad R_t^{\theta} \geq R_*.$$

It is in the carrier-thin phase if

$$\rho_t^{\mathbb{Q}} \geq \rho_*, \quad K_t^{\text{car}} < K_*.$$

When regions overlap, residual-stress and carrier-thin labels take precedence because they mark failure of public carrier absorption or enlargement of the residual price envelope.

The thresholds in Definition 5.3 are not fitted trading cutoffs. They are coordinate surfaces on the recovered $\mathbb{Q}$ -state. The spectral surface separates local and synchronized option geometry; the carrier surface separates synchronized risk absorbed by the admitted carrier span from synchronized risk that remains thinly spanned; the residual surface marks expansion of public-equivalent price freedom.

The intrinsic phase atlas precedes any positive finite-resolution band. Its quotient seams are the surfaces where an observability or pricing mechanism changes type: normalized Perron amplification crosses its critical level, the Perron spectral gap collapses, the carrier Gram lower bound reaches zero, a maturity or route-rank determinant vanishes, or the residual price envelope collapses. Positive constants (ρ_*, K_*, R_*) thicken these seams into certificate regions for finite panels. Changing the bands changes the resolution of the certificate, not the underlying quotient stratification.

Theorem 5.4 (Risk-neutral phase stratification). *On a recovery stratum, the phase labels in Definition 5.3 are gauge-invariant functions of the recovered source quotient. Their transition surfaces are quotient-invariant rank, gap, Gram, and residual-width surfaces:*

$$\rho^{\mathbb{Q}} = \rho_*, \quad \lambda_{\min}(\Gamma_{\text{car}}) = K_*, \quad R^{\theta} = R_*.$$

If the recovered state process is adapted and continuous, phase changes can occur only by hitting one of these surfaces. Thus local, synchronized-carrier, residual-stress, and carrier-thin phases are not additional latent regimes; they are strata of the recovered risk-neutral option quotient.

Proof. The coordinates $\rho^\mathbb{Q}$, Γ_{car} , and R^θ are defined from the canonical recovered state and the selected residual rule, so they are invariant under the gauge normalizations used in the source quotient. The phase regions are inverse images of Borel subsets of these quotient coordinates. If the recovered state is adapted and continuous, the phase coordinates are adapted and continuous. A path cannot move between open phase regions without crossing the boundary surface that separates them. $\square$

Remark 5.5 (Spectral-gap seams). The surface $\Delta_P = 0$, where the Perron spectral gap closes, should be read as a spectral discriminant seam rather than as an ordinary smooth phase boundary. On finite self-adjoint or symmetry-reduced subfamilies such seams may appear as familiar eigenvalue-intersection loci; in non-self-adjoint directed Perron operators their codimension depends on the admitted source family, positivity constraints, and whether the crossing is an eigenvalue collision or a dominance exchange in modulus. The theory uses only the invariant fact that the reduced resolvent and the differential of the Perron projector lose uniform bounds as $\Delta_P \downarrow 0$, and may also lose bounds when non-normal conditioning degenerates; no universal codimension assignment is required. Thus a path hitting this seam exits the smooth recovered phase chart and must be continued, if possible, in another chart of the recovery atlas.

Definition 5.6 (Smooth phase weights). For differentiable pricing, replace the hard phase labels by smooth weights. Define scores

$$s_\rho = \rho_t^\mathbb{Q} - \rho_*, \quad s_K = K_t^{\text{car}} - K_*, \quad s_R = R_t^\theta - R_*, \quad s_D = D_t^\mathbb{Q} - D_*,$$

and

$$\begin{aligned} m_{\text{loc}} &= -s_\rho - s_R, & m_{\text{syn}} &= s_\rho + s_K - s_R + \eta_D s_D, \\ m_{\text{res}} &= s_\rho + s_R + \eta_D s_D, & m_{\text{thin}} &= s_\rho - s_K. \end{aligned}$$

Here D_* is a directed-pressure reference level and $\eta_D \geq 0$ controls how strongly directed route pressure tilts synchronized and residual phases. For a sharpness parameter $\rho > 0$, set

$$\pi_{\nu,t} = \frac{\exp(\rho m_{\nu,t})}{\sum_{\mu \in \{\text{loc}, \text{syn}, \text{res}, \text{thin}\}} \exp(\rho m_{\mu,t})}, \quad \nu \in \{\text{loc}, \text{syn}, \text{res}, \text{thin}\}.$$

The vector π_t is the smooth risk-neutral phase allocation induced by the recovered Volterra–Perron state.

Definition 5.7 (Phase-weighted residual selector). For each $\nu \in \{\text{loc}, \text{syn}, \text{res}, \text{thin}\}$, let

$$Z_{t,S}^\nu = 1 + \sum_r \theta_{r,t}^\nu R^r$$

be a positive public-equivalent density family, indexed by maturity $S > t$. The phase-weighted family is

$$Z_{t,S}^\pi := \sum_\nu \pi_{\nu,t} Z_{t,S}^\nu.$$

The corresponding selected price is

$$\Pi_t^\pi(X) := \mathbb{E}^\mathbb{Q}[Z_{t,T}^\pi X \mid \mathcal{H}_t].$$

Proposition 5.8 (Phase transitions are adapted). *Suppose the recovered state process $\Xi_t^\mathbb{Q}$, together with the selected residual-density rule θ_t , is $\mathcal{H}_t$ -adapted and the phase coordinates in Definition 5.1 are continuous functions of the canonical state. Then the hard phase label is $\mathcal{H}_t$ -adapted, and the first hitting times of the phase transition surfaces*

$$\{\rho^\mathbb{Q} = \rho_*\}, \quad \{K^{\text{car}} = K_*\}, \quad \{R^\theta = R_*\}$$

are $(\mathcal{H}_t)$ -stopping times. If the recovered state is continuous, phase changes occur only by hitting one of these transition surfaces.

Proof. The phase coordinates are adapted as continuous functions of the adapted recovered state. The hard phase regions are Borel sets in the canonical state chart, so the phase label is adapted. The transition surfaces are closed level sets of adapted continuous processes. Their first hitting times are therefore stopping times. If the recovered state path is continuous, the phase label cannot move from one open region to another without crossing the boundary separating them. $\square$

Proposition 5.9 (No-arbitrage preservation under phase weighting). *If each $Z_{t,S}^\nu$ in Definition 5.7 is positive and public-equivalent,*

$$\mathbb{E}^\mathbb{Q}[Z_{t,S}^\nu \mid \mathcal{H}_t] = 1, \quad S > t,$$

and

$$\mathbb{E}^\mathbb{Q}[Z_{t,T(a)}^\nu X_a \mid \mathcal{H}_t] = \mathbb{E}^\mathbb{Q}[X_a \mid \mathcal{H}_t] \quad \text{for every admitted public payoff } X_a,$$

and the weights $\pi_{\nu,t}$ are nonnegative, $\mathcal{H}_t$ -measurable, and sum to one, then

$$Z_{t,S}^\pi = \sum_\nu \pi_{\nu,t} Z_{t,S}^\nu$$

is positive and public-equivalent. Hence phase-weighted residual pricing preserves all public option prices while allowing the latent residual charge to vary with the recovered market phase.

Proof. Positivity follows from positivity of the kernels and nonnegativity of the weights. Since the weights are $\mathcal{H}_t$ -measurable,

$$\mathbb{E}^\mathbb{Q}[Z_{t,S}^\pi \mid \mathcal{H}_t] = \sum_\nu \pi_{\nu,t} \mathbb{E}^\mathbb{Q}[Z_{t,S}^\nu \mid \mathcal{H}_t] = \sum_\nu \pi_{\nu,t} = 1.$$

For any admitted public payoff X_a ,

$$\mathbb{E}^\mathbb{Q}[Z_{t,T(a)}^\pi X_a \mid \mathcal{H}_t] = \sum_\nu \pi_{\nu,t} \mathbb{E}^\mathbb{Q}[Z_{t,T(a)}^\nu X_a \mid \mathcal{H}_t] = \mathbb{E}^\mathbb{Q}[X_a \mid \mathcal{H}_t].$$

Thus the kernel is public-equivalent. Public payoff prices are unchanged, while residual payoffs can receive phase-dependent charges through their covariance with the selected residual directions. $\square$

Remark 5.10 (Datewise preservation versus process consistency). Proposition 5.9 is a datewise preservation result. It says that, at a fixed date and horizon, convex phase weighting preserves the admitted public payoff prices. It does not by itself produce a single density process whose horizon ratios generate all the $Z_{t,S}^\pi$ consistently across t and S . Dynamic no-arbitrage requires the stronger process-consistent kernel of Definition 10.10 and the martingale restrictions of Section 10. Thus phase weighting is a public residual-selection rule until it is embedded in a martingale density process.

Proposition 5.11 (Semimartingale phase dynamics). *Suppose the finite-dimensional recovered chart U_t satisfies the semimartingale dynamics in Assumption 10.3, and suppose the score maps*

$$U \mapsto m_\nu(U), \quad \nu \in \{\text{loc}, \text{syn}, \text{res}, \text{thin}\},$$

are C^2 on the recovery neighborhood. Then the smooth phase weights $\pi_{\nu,t}$ are $\mathcal{H}_t$ -semimartingales. Their drift and diffusion coefficients are given by Itô's formula applied to the softmax map in Definition 5.6.

Proof. The recovered chart is an $\mathcal{H}_t$ -semimartingale. The score vector is a C^2 function of that chart, and the softmax map is C^∞ . The composition is therefore an $\mathcal{H}_t$ -semimartingale, with coefficients given by Itô's formula. $\square$

Definition 5.12 (Phase value functions). Let U_t be a finite-dimensional recovered state chart that is Markov for the selected claim family under the phase kernels below. For a claim X and phase $\nu \in \{\text{loc}, \text{syn}, \text{res}, \text{thin}\}$, define a regular conditional version

$$\Phi_\nu(t, u; X) := \mathbb{E}^\mathbb{Q}[Z_{t,T}^\nu X \mid U_t = u], \quad \Phi_\pi(t, u; X) := \sum_\nu \pi_\nu(u) \Phi_\nu(t, u; X).$$

The phase transition charge relative to a reference phase ν_0 is

$$\mathfrak{C}_{\nu_0}^{\text{ph}}(t, u; X) := \Phi_\pi(t, u; X) - \Phi_{\nu_0}(t, u; X) = \sum_{\nu} \pi_\nu(u) (\Phi_\nu(t, u; X) - \Phi_{\nu_0}(t, u; X)).$$

Definition 5.13 (Transition Greeks). Let ξ be one of the phase scores s_ρ, s_K, s_R . The pure allocation transition Greek of X in the ξ -direction is

$$\mathfrak{T}_\xi^{\text{alloc}} X := \sum_{\nu} \frac{\partial \pi_\nu}{\partial \xi} \Phi_\nu(t, U_t; X).$$

Using the softmax weights of Definition 5.6,

$$\mathfrak{T}_\xi^{\text{alloc}} X = \rho \sum_{\nu} \pi_{\nu,t} \left(\partial_\xi m_{\nu,t} - \sum_{\mu} \pi_{\mu,t} \partial_\xi m_{\mu,t} \right) \Phi_\nu(t, U_t; X).$$

The full transition Greek adds the direct state sensitivity

$$\mathfrak{T}_\xi X := \mathfrak{T}_\xi^{\text{alloc}} X + \sum_{\nu} \pi_{\nu,t} \partial_\xi \Phi_\nu(t, U_t; X),$$

whenever the displayed derivatives exist.

Theorem 5.14 (Phase-weighted risk-neutral pricing equation). *Suppose the recovered chart satisfies*

$$dU_t = b(t, U_t) dt + \Sigma(t, U_t) dW_t^{\mathbb{Q}}$$

on a recovery neighborhood, and define

$$\mathcal{A}f := Df[b] + \frac{1}{2} \text{Tr}(\Sigma^* D^2 f \Sigma).$$

Let the phase weights π_ν and phase value functions Φ_ν be $C^{1,2}$. If the discounted phase-weighted selected price is a local $\mathbb{Q}$ -martingale, then

$$(\partial_t + \mathcal{A} - r_t) \Phi_\pi(t, U_t; X) = 0.$$

Equivalently,

$$0 = \sum_{\nu} \pi_\nu (\partial_t + \mathcal{A} - r_t) \Phi_\nu + \sum_{\nu} \Phi_\nu (\partial_t + \mathcal{A}) \pi_\nu + \sum_{\nu} \langle \Sigma^* D \pi_\nu, \Sigma^* D \Phi_\nu \rangle,$$

evaluated at (t, U_t) . The last two sums are the transition drift and transition covariance charges generated by phase weights that move with the recovered $\mathbb{Q}$ -state.

Proof. The first display is Itô's formula applied to $e^{-\int_0^t r_s ds} \Phi_\pi(t, U_t; X)$ and the local martingale condition. For the expanded form, use

$$\Phi_\pi = \sum_{\nu} \pi_\nu \Phi_\nu$$

and apply the product rule to $(\partial_t + \mathcal{A})(\pi_\nu \Phi_\nu)$. The second-order cross term is

$$\langle \Sigma^* D \pi_\nu, \Sigma^* D \Phi_\nu \rangle,$$

which gives the displayed transition covariance charge. $\square$

Corollary 5.15 (Phase charges vanish only under frozen or balanced phases). *If every phase value function separately satisfies*

$$(\partial_t + \mathcal{A} - r_t) \Phi_\nu = 0,$$

then the phase-weighted selected price is arbitrage-free only if

$$\sum_{\nu} \Phi_\nu (\partial_t + \mathcal{A}) \pi_\nu + \sum_{\nu} \langle \Sigma^* D \pi_\nu, \Sigma^* D \Phi_\nu \rangle = 0.$$

In particular, changing risk-neutral phase weights are priced objects unless their drift and diffusion effects are exactly balanced by phase-value spreads.

Proof. Substitute the separate phase pricing equations into Theorem 5.14. The remaining terms must vanish for the phase-weighted discounted price to be a local martingale. $\square$

6. PUBLIC NON-IDENTIFICATION

Public compression gives the first obstruction. Full latent recovery requires a source-family restriction before the public quotient can be inverted.

Definition 6.1 (Public indistinguishability). Fix date t and an admitted option-pricing map

$$\mathfrak{P}_t : \mathfrak{X}_t \rightarrow \mathcal{Y}_t, \quad \Theta \mapsto \mathcal{M}_t^{\text{opt}}(\Theta).$$

Two latent states are public-indistinguishable, written

$$\Theta \sim_{\text{pub},t} \Theta',$$

if $\mathfrak{P}_t(\Theta) = \mathfrak{P}_t(\Theta')$ on every admitted public option coordinate at date t . The public-observable state space is the quotient

$$\mathfrak{X}_t^{\text{pub}} := \mathfrak{X}_t / \sim_{\text{pub},t},$$

with canonical projection $\pi_t^{\text{pub}} : \mathfrak{X}_t \rightarrow \mathfrak{X}_t^{\text{pub}}$.

Theorem 6.2 (Maximality of the public-observable quotient). *Let $Z : \mathfrak{X}_t \rightarrow \mathcal{Z}$ be any state functional. The following are equivalent.*

- (1) Z is identified by the public option manifold, in the sense that there is a map $\zeta : \mathfrak{P}_t(\mathfrak{X}_t) \rightarrow \mathcal{Z}$ such that $Z = \zeta \circ \mathfrak{P}_t$.
- (2) Z is constant on public-indistinguishability classes.
- (3) Z factors through the quotient, in the sense that there is a map $\bar{Z} : \mathfrak{X}_t^{\text{pub}} \rightarrow \mathcal{Z}$ such that

$$Z = \bar{Z} \circ \pi_t^{\text{pub}}.$$

Consequently, $\mathfrak{X}_t^{\text{pub}}$ is the largest latent object identifiable from the admitted public option manifold without additional source-family restrictions.

Proof. If $Z = \zeta \circ \mathfrak{P}_t$, equal public manifolds imply equal Z -values, so Z is constant on public-indistinguishability classes. Conversely, if Z is constant on those classes, define

$$\bar{Z}([\Theta]_{\text{pub}}) := Z(\Theta).$$

The definition is independent of the representative exactly because the classes are the fibers of $\mathfrak{P}_t$. It also defines a map on the public image by

$$\zeta(\mathfrak{P}_t(\Theta)) := Z(\Theta),$$

again well defined on each fiber. Thus $Z = \bar{Z} \circ \pi_t^{\text{pub}} = \zeta \circ \mathfrak{P}_t$. Every identified functional factors through $\mathfrak{X}_t^{\text{pub}}$, which is the maximal public-observable quotient. $\square$

Corollary 6.3 (Source-family recovery as quotient embedding). *Let $\mathfrak{S}_t \subset \mathfrak{X}_t$ be a gauge-normalized Volterra–Perron source family, and let $\mathfrak{S}_t / \sim_g$ be its gauge quotient. Local recovery from the public manifold is equivalent to local injectivity of*

$$\mathfrak{S}_t / \sim_g \longrightarrow \mathfrak{X}_t^{\text{pub}}, \quad [\Theta]_g \mapsto [\Theta]_{\text{pub}}.$$

Thus the source-family rank conditions are exactly the conditions under which the canonical source quotient embeds into the public-observable quotient.

Proof. By Theorem 6.2, the public manifold identifies exactly the public quotient $\mathfrak{X}_t^{\text{pub}}$. A recovery map on the source family must therefore invert the induced map from the gauge quotient $\mathfrak{S}_t / \sim_g$ into this public quotient. Local recovery is local injectivity of that induced map, and stable local recovery is the same statement with a positive inverse modulus. The source-family rank conditions are the differential certificate for this quotient embedding. $\square$

Definition 6.4 (Dynamic public-equivalent density family). At date t , a positive density family

$$Z_{t,\cdot} = (Z_{t,S})_{S>t}, \quad Z_{t,S} \in L^1(\mathbb{Q}, \mathcal{K}_S),$$

is public-equivalent if, for every $S > t$,

$$\mathbb{E}^{\mathbb{Q}}[Z_{t,S} \mid \mathcal{H}_t] = 1$$

and, for every admitted public payoff X_a with maturity $T(a)$,

$$\mathbb{E}^{\mathbb{Q}}[Z_{t,T(a)}X_a \mid \mathcal{H}_t] = \mathbb{E}^{\mathbb{Q}}[X_a \mid \mathcal{H}_t].$$

It prices a claim $X \in L^1(\mathbb{Q}, \mathcal{F}_T)$ by

$$\Pi_t^Z(X) := \mathbb{E}^{\mathbb{Q}}[Z_{t,T}X \mid \mathcal{H}_t].$$

Lemma 6.5 (Public preservation). *If $Z_{t,\cdot}$ is public-equivalent and $Y \in L^\infty(\mathbb{Q}, \mathcal{H}_t)$, then, for every $S > t$,*

$$\mathbb{E}^{\mathbb{Q}}[Z_{t,S}Y \mid \mathcal{H}_t] = Y.$$

Consequently the density family preserves every date- t public payoff test.

Proof. The public-equivalence normalization is $\mathbb{E}^{\mathbb{Q}}[Z_{t,S} \mid \mathcal{H}_t] = 1$. Since Y is $\mathcal{H}_t$ -measurable,

$$\mathbb{E}^{\mathbb{Q}}[Z_{t,S}Y \mid \mathcal{H}_t] = Y \mathbb{E}^{\mathbb{Q}}[Z_{t,S} \mid \mathcal{H}_t] = Y.$$

Multiplication by a date- t public test therefore does not leave the public equivalence class. $\square$

Lemma 6.6 (Residual annihilators generate public-equivalent density families). *Let $R \in L^\infty(\mathbb{Q}, \mathcal{K}_t)$ satisfy*

$$\mathbb{E}^{\mathbb{Q}}[R \mid \mathcal{H}_t] = 0, \quad \mathbb{E}^{\mathbb{Q}}[RX_a \mid \mathcal{H}_t] = 0 \quad \text{for every admitted public payoff } X_a.$$

If $|\eta| < \|R\|_\infty^{-1}$, then the density family

$$Z_{t,S}^\eta = 1 + \eta R, \quad S > t,$$

is positive and public-equivalent.

Proof. The bound $|\eta| < \|R\|_\infty^{-1}$ gives $1 + \eta R \geq 1 - |\eta| \|R\|_\infty > 0$ almost surely. The normalization condition gives

$$\mathbb{E}^{\mathbb{Q}}[Z_{t,S}^\eta \mid \mathcal{H}_t] = 1 + \eta \mathbb{E}^{\mathbb{Q}}[R \mid \mathcal{H}_t] = 1.$$

For every admitted payoff X_a , the annihilator condition gives

$$\mathbb{E}^{\mathbb{Q}}[Z_{t,T(a)}^\eta X_a \mid \mathcal{H}_t] = \mathbb{E}^{\mathbb{Q}}[X_a \mid \mathcal{H}_t] + \eta \mathbb{E}^{\mathbb{Q}}[RX_a \mid \mathcal{H}_t] = \mathbb{E}^{\mathbb{Q}}[X_a \mid \mathcal{H}_t].$$

The density family is positive and public-equivalent. $\square$

Theorem 6.7 (Public non-identification of latent stress). *Fix $t < T$. Suppose there is a bounded $R \in L^\infty(\mathbb{Q}, \mathcal{K}_t)$ such that*

$$(A) \quad \mathbb{E}^{\mathbb{Q}}[R \mid \mathcal{H}_t] = 0 \quad \text{and} \quad \mathbb{E}^{\mathbb{Q}}[RX_a \mid \mathcal{H}_t] = 0 \quad \text{for every admitted public payoff } X_a,$$

but

$$\mathbb{E}^{\mathbb{Q}}[L_{t,T}^{\mathbb{Q}}(X)R \mid \mathcal{H}_t] \neq 0$$

on a set of positive probability for some latent-stress claim X . Then, for sufficiently small nonzero η , the density family

$$Z_{t,S}^\eta = 1 + \eta R, \quad S > t,$$

preserves the public option manifold at date t but changes the selected price of X :

$$\Pi_t^{Z^\eta}(X) = \mathbb{E}^{\mathbb{Q}}[X \mid \mathcal{H}_t] + \eta \mathbb{E}^{\mathbb{Q}}[L_{t,T}^{\mathbb{Q}}(X)R \mid \mathcal{H}_t].$$

Thus public calibration alone cannot recover the latent stress price.

Proof. Choose $|\eta| < \|R\|_\infty^{-1}$. Lemma 6.6 shows that $Z_{t,\cdot}^\eta$ is public-equivalent and therefore preserves the public option manifold. For the latent claim,

$$\mathbb{E}^{\mathbb{Q}}[Z_{t,T}^\eta X \mid \mathcal{H}_t] = \mathbb{E}^{\mathbb{Q}}[X \mid \mathcal{H}_t] + \eta \mathbb{E}^{\mathbb{Q}}[XR \mid \mathcal{H}_t].$$

Since $R \in L^\infty(\mathbb{Q}, \mathcal{K}_t)$, the tower property gives

$$\mathbb{E}^{\mathbb{Q}}[XR \mid \mathcal{H}_t] = \mathbb{E}^{\mathbb{Q}}[\mathbb{E}^{\mathbb{Q}}[X \mid \mathcal{K}_t]R \mid \mathcal{H}_t] = \mathbb{E}^{\mathbb{Q}}[L_{t,T}^{\mathbb{Q}}(X)R \mid \mathcal{H}_t].$$

The last term is nonzero on a set of positive probability by assumption, so the selected price is moved while all public prices are fixed. $\square$

Corollary 6.8 (No universal recovery from public compression). *If the residual space $L^2(\mathbb{Q}, \mathcal{K}_t) \ominus L^2(\mathbb{Q}, \mathcal{H}_t)$ contains a bounded direction correlated with a latent-stress loading, then no procedure using only the public option manifold can identify all latent-stress prices without additional source-family restrictions.*

Proof. Theorem 6.7 constructs two positive public-equivalent density families with the same value of every admitted public payoff and different conditional prices for the latent claim. A reconstruction rule whose input is only $\mathcal{M}_t^{\text{opt}}$ has the same input under both families. It therefore cannot identify both selected latent-stress prices. $\square$

Corollary 6.9 (Exact residual obstruction). *Fix the admitted public payoff family at date t , and write*

$$L_a := \mathbb{E}^{\mathbb{Q}}[X_a \mid \mathcal{K}_t], \quad L_X := \mathbb{E}^{\mathbb{Q}}[X \mid \mathcal{K}_t].$$

Let

$$\mathcal{P}_t^{\text{pub}} := \overline{\left\{ \eta_0 + \sum_{j=1}^n \eta_j L_{a_j} : n < \infty, a_j \in \mathcal{I}_t, \eta_0, \eta_j \in L^\infty(\mathbb{Q}, \mathcal{H}_t) \right\}}^{L^2(\mathbb{Q}, \mathcal{K}_t)} \subset L^2(\mathbb{Q}, \mathcal{K}_t)$$

be the closed conditional public-loading module generated by the date- t public payoffs. Assume bounded elements are dense in $(\mathcal{P}_t^{\text{pub}})^\perp$. A latent claim X has a public-equivalent residual price degree of freedom at date t if and only if

$$(I - C_{\mathcal{P}_t^{\text{pub}}}^{\mathbb{Q}})L_X \neq 0.$$

In that case there is a bounded residual direction R preserving all admitted public prices and changing the selected price of X .

Proof. Orthogonality to the conditional module is equivalent to the conditional annihilator equations:

$$R \in (\mathcal{P}_t^{\text{pub}})^\perp \iff \mathbb{E}^{\mathbb{Q}}[R \mid \mathcal{H}_t] = 0, \quad \mathbb{E}^{\mathbb{Q}}[RL_a \mid \mathcal{H}_t] = 0 \quad \text{for every } a \in \mathcal{I}_t.$$

The equivalence follows by testing R against $Y(\eta_0 + \sum_j \eta_j L_{a_j})$ with bounded $\mathcal{H}_t$ -measurable coefficients and using the defining property of conditional expectation. Since R is $\mathcal{K}_t$ -measurable, $\mathbb{E}^{\mathbb{Q}}[RX_a \mid \mathcal{H}_t] = \mathbb{E}^{\mathbb{Q}}[RL_a \mid \mathcal{H}_t]$, so these are exactly the public-equivalence annihilator equations.

If $(I - C_{\mathcal{P}_t^{\text{pub}}}^{\mathbb{Q}})L_X = 0$, then L_X belongs to the conditional public-loading module, and every public-annihilating residual direction has $\mathbb{E}^{\mathbb{Q}}[L_X R \mid \mathcal{H}_t] = 0$. Public-equivalent tilts cannot change the price of X .

If the projection is nonzero, set

$$R_0 = (I - C_{\mathcal{P}_t^{\text{pub}}}^{\mathbb{Q}})L_X.$$

Then $R_0 \in (\mathcal{P}_t^{\text{pub}})^\perp$ and $\mathbb{E}^{\mathbb{Q}}[L_X R_0 \mid \mathcal{H}_t] = \mathbb{E}^{\mathbb{Q}}[R_0^2 \mid \mathcal{H}_t]$ is nonzero on the event where the residual projection is nonzero. By the density assumption, choose a bounded residual direction $R \in (\mathcal{P}_t^{\text{pub}})^\perp$ sufficiently close to R_0 in conditional L^2 to preserve a nonzero conditional pairing with L_X . Lemma 6.6 then generates positive public-equivalent density families that leave the public manifold fixed and move the price of X . $\square$

Definition 6.10 (Public payoff-module completion). For a fixed horizon T , define the closed conditional public payoff module

$$\mathcal{A}_{t,T}^{\text{pub}} := \overline{\text{span}\{Y L_{t,T(a)}^{\mathbb{Q}}(X_a) : Y \in L^\infty(\mathbb{Q}, \mathcal{H}_t), a \in \mathcal{I}_t, T(a) = T\}}^{L^2(\mathbb{Q}, \mathcal{K}_t)} \subset L^2(\mathbb{Q}, \mathcal{K}_t).$$

The admitted public manifold is complete for the date- t latent loading class $\mathcal{L}_{t,T} \subset L^2(\mathbb{Q}, \mathcal{K}_t)$ if

$$\mathcal{L}_{t,T} \subseteq \mathcal{A}_{t,T}^{\text{pub}}.$$

Theorem 6.11 (Public completion and residual collapse). *Fix $t < T$, and let $\mathcal{L}_{t,T} \subset L^2(\mathbb{Q}, \mathcal{K}_t)$ be the latent loading class whose claims are to be priced. The following are equivalent.*

- (1) The public manifold is complete for $\mathcal{L}_{t,T}$.
- (2) Every $L \in \mathcal{L}_{t,T}$ has zero orthogonal residual relative to $\mathcal{A}_{t,T}^{\text{pub}}$.
- (3) No public-equivalent density perturbation generated by a bounded public-annihilating residual direction can change the price of any claim with loading in $\mathcal{L}_{t,T}$.

When these conditions hold, residual price intervals collapse to singletons on $\mathcal{L}_{t,T}$. If they fail and bounded annihilating residual directions are dense in the orthogonal complement, every nonzero loading residual generates a genuine residual pricing fiber.

Proof. The first two statements are the Hilbert projection decomposition of each loading $L \in \mathcal{L}_{t,T}$ into its component in $\mathcal{A}_{t,T}^{\text{pub}}$ and its orthogonal residual. If every residual component is zero, then every public-annihilating direction $R \in (\mathcal{A}_{t,T}^{\text{pub}})^\perp$ has zero conditional pairing with every such loading. Lemma 6.6 therefore cannot move the selected price of any claim in the class.

Conversely, if some loading has a nonzero orthogonal residual, the density assumption supplies a bounded annihilating direction with nonzero conditional pairing against that residual. The positive densities generated by Lemma 6.6 preserve all admitted public prices and move the claim price. The image of this one-dimensional tilt is a nontrivial residual pricing fiber, and it collapses exactly when the projection residual is zero. $\square$

Corollary 6.12 (Selector redundancy at public completion). *On a publicly complete latent loading class, the residual selector θ_t is redundant. Outside public completion, the selector chooses a point in the unspanned pricing fiber left by the public payoff module.*

Proof. Theorem 6.11 identifies public completion with vanishing residual projection for every loading in the claim class. In that case all public-annihilating density tilts have zero conditional pairing with those loadings, so the selected residual term is fixed and the selector carries no pricing degree of freedom. If completion fails, the orthogonal residual projection is nonzero for some loading; Lemma 6.6 then gives public-equivalent tilts that move the corresponding price. The selector chooses one element of that residual fiber. $\square$

Theorem 6.13 (Option-manifold impossibility theorem). *Fix date t and an arbitrage-free admitted public option family. The following obstructions are sharp.*

- (1) Without source-family restrictions, the public manifold identifies only $\mathfrak{X}_t^{\text{pub}}$. No calibration or estimator using only the admitted public coordinates can distinguish two latent states in the same public equivalence class.
- (2) Single-name vanilla surfaces identify marginal risk-neutral laws, not joint covariance routes, directed Volterra contagion, or Perron synchronization. Distinct couplings of the same marginals may have the same single-name surfaces and different index or latent-stress prices.
- (3) If the finite basket or carrier design has a nontrivial null direction on the active covariance-route subspace, then option prices on that finite panel cannot separate perturbations in that null direction while they remain inside the covariance cone.
- (4) If the Perron reduced-resolvent norm diverges, in particular along gap-closing seams, the Perron projector is not uniformly Lipschitz as a function of the public response; any recovery map for $e^{\mathbb{Q}}$, $\ell^{\mathbb{Q}}$, or $P^{\mathbb{Q}}$ has an inverse modulus that diverges along such strata.
- (5) If the maturity dictionary has a nontrivial null combination on the observed maturities, distinct Volterra source geometries produce the same maturity panel.
- (6) If the public payoff module is incomplete for a claim class, public-equivalent residual density families preserve every admitted public option price while assigning different prices to some latent-stress claims in that class.

Thus every positive recovery statement requires the corresponding rank, gap, source-family, or public-completion hypothesis. The failures are mathematical obstructions in the option manifold, not deficiencies of a particular numerical procedure.

Proof. The first statement is Theorem 6.2. For the second, fixed single-name marginal laws do not determine a joint coupling; distinct couplings with the same marginals can change cross-covariances and hence index or latent-stress prices. The third is the nullspace argument

used in the covariance-recovery and elliptope results of Theorems 7.11 and 7.12. The fourth follows from the reduced-resolvent formula for a simple isolated eigenvalue, because the derivative of the Perron projector is controlled by that resolvent. The fifth is rank failure of the finite maturity response matrix. The sixth is Theorem 6.11 together with Lemma 6.6. Each obstruction therefore leaves either a kernel in the public response, an unbounded inverse modulus, or a nontrivial residual pricing fiber. $\square$

7. RISK-NEUTRAL OBSERVABILITY OF THE OPTION MANIFOLD

Recovery requires more than public calibration. The option manifold must separate canonical coordinates after gauge directions have been removed. This section gives the infinitesimal and finite-panel verification for Theorem 8.22. Observability is the derivative-level form of quotient embedding; the smallest quotient singular value is the local conditioning scale for source-chart recovery from public option data.

Definition 7.1 (Risk-neutral observability operator). Let $\psi : U \subset E \rightarrow \mathfrak{X}_t^{\text{str}}$ be a gauge-normalized source chart and let

$$F = \mathcal{R}_t \circ \psi : U \rightarrow \mathcal{Y}_t$$

be the observable option-response map. The risk-neutral observability operator at u_0 is

$$\mathcal{O}_t^{\mathbb{Q}} := DF(u_0) : E \rightarrow \mathcal{Y}_t.$$

Its canonical version is the induced quotient operator

$$\overline{\mathcal{O}}_t^{\mathbb{Q}} : \bar{E} = E/E_g \rightarrow \text{Ran } DF(u_0), \quad \overline{\mathcal{O}}_t^{\mathbb{Q}}[h] = DF(u_0)h.$$

The state is locally observable at u_0 if $\overline{\mathcal{O}}_t^{\mathbb{Q}}$ has a bounded left inverse on its range.

Definition 7.2 (Single-name and joint option manifolds). The single-name vanilla manifold is the restriction of $\mathcal{M}_t^{\text{opt}}$ to payoffs depending on one underlier at a time. The joint index/carrier manifold adds index, basket, variance, or carrier-spread option coordinates whose payoffs depend on at least two underliers through the same maturity window.

Theorem 7.3 (Vanilla insufficiency for joint stress). *Suppose two risk-neutral Volterra–Perron states have the same conditional marginal law for each single underlier over all maturities in the admitted vanilla family, but have different conditional cross-variation response*

$$\text{Cov}^{\mathbb{Q}}\left(\int_t^{t+\tau} V_i(s) ds, \int_t^{t+\tau} V_j(s) ds \mid \mathcal{H}_t\right)$$

for some $i \neq j$ and some maturity τ . Then their single-name vanilla manifolds agree, but a nondegenerate index or basket variance response differs. Consequently, single-name vanilla surfaces alone do not recover the directed Volterra contagion operator, the Perron synchronization state, or synchronized latent stress.

Proof. A single-name vanilla payoff is a Borel function of one terminal underlier. Its date- t price is determined by the corresponding conditional marginal law. Equal conditional marginal laws therefore imply identical single-name vanilla surfaces. An index or basket variance claim contains cross terms of the form

$$2\omega_i\omega_j \text{Cov}^{\mathbb{Q}}\left(\int_t^{t+\tau} V_i(s) ds, \int_t^{t+\tau} V_j(s) ds \mid \mathcal{H}_t\right),$$

for nonzero weights $\omega_i\omega_j$. If this cross response differs and the basket denominator is nondegenerate, the joint option response differs. The directed Volterra and Perron synchronization coordinates enter exactly through these joint responses, so they are not identified by marginal vanilla prices alone. $\square$

Proposition 7.4 (Basket design observability). *Let $\mathcal{E}$ be the active set of unordered cross-pairs (i, j) used in recovery. For each admitted basket or index u , define the cross-variance response*

$$B_u(\tau) := \text{TV}_{u,t}(\tau) - \text{TV}_{u,t}^0(\tau) - \sum_i (w_i^u)^2 V_{ii,t}^{\mathbb{Q}}(\tau) = 2 \sum_{(i,j) \in \mathcal{E}} w_i^u w_j^u V_{ij,t}^{\mathbb{Q}}(\tau).$$

Let the basket design matrix be

$$W_{u,(i,j)} := 2w_i^u w_j^u, \quad (i,j) \in \mathcal{E}.$$

Then, for each maturity τ , the active cross-route vector

$$V_{\mathcal{E}}(\tau) = (V_{ij,t}^{\mathbb{Q}}(\tau))_{(i,j) \in \mathcal{E}}$$

is identified from the joint basket responses $B(\tau)$ if and only if

$$\text{rank } W_{\mathcal{E}} = |\mathcal{E}|.$$

When this rank condition holds,

$$\|\delta V_{\mathcal{E}}(\tau)\| \leq \sigma_{\min}(W_{\mathcal{E}})^{-1} \|\delta B(\tau)\|.$$

Proof. The displayed variance decomposition gives the linear system

$$B(\tau) = W_{\mathcal{E}} V_{\mathcal{E}}(\tau).$$

The vector $V_{\mathcal{E}}(\tau)$ is identifiable exactly when this linear map is injective on the active cross-route coordinates, which is equivalent to full column rank. If $W_{\mathcal{E}}$ has full column rank, its Moore–Penrose left inverse has norm $\sigma_{\min}(W_{\mathcal{E}})^{-1}$, which gives the displayed stability estimate. $\square$

Definition 7.5 (Dual recovery certificate). Let $F : U \subset \mathbb{R}^d \rightarrow \mathbb{R}^m$ be a finite quote response map with full-column-rank Jacobian

$$J := DF(u_0).$$

A dual recovery certificate for latent coordinate k is a vector $a_k \in \mathbb{R}^m$ satisfying

$$a_k^{\top} J = e_k^{\top},$$

where e_k is the k th coordinate vector in $\mathbb{R}^d$. A matrix $A \in \mathbb{R}^{d \times m}$ with $AJ = I_d$ is a full dual certificate.

Theorem 7.6 (Quote portfolios recover infinitesimal coordinates). Let $F : U \subset \mathbb{R}^d \rightarrow \mathbb{R}^m$ be C^2 , and suppose $J = DF(u_0)$ has full column rank. If A is any left inverse of J , then, for u near u_0 ,

$$A(F(u) - F(u_0)) = u - u_0 + O(\|u - u_0\|^2).$$

Equivalently,

$$u_k - u_{0,k} = a_k^{\top} (F(u) - F(u_0)) + O(\|u - u_0\|^2).$$

If quote noise has covariance matrix Σ_y , the linear certificate variance for coordinate k is $a_k^{\top} \Sigma_y a_k$. When Σ_y is positive definite, the minimum-variance linear unbiased certificate is the k th row of

$$A_{\Sigma} := (J^{\top} \Sigma_y^{-1} J)^{-1} J^{\top} \Sigma_y^{-1}.$$

Proof. Taylor expansion gives $F(u) - F(u_0) = J(u - u_0) + O(\|u - u_0\|^2)$. Multiplication by A and the identity $AJ = I_d$ give the coordinate recovery expansion. A certificate a_k applies the linear quote functional $a_k^{\top} \delta y$, so its noise variance is $a_k^{\top} \Sigma_y a_k$. Minimizing this variance subject to the unbiasedness constraint $AJ = I_d$ gives the normal equations $A \Sigma_y^{-1} J = I_d$, and therefore the generalized least-squares left inverse $(J^{\top} \Sigma_y^{-1} J)^{-1} J^{\top} \Sigma_y^{-1}$. $\square$

Corollary 7.7 (Functional certificate in infinite quote spaces). Let $DF(u_0) : E \rightarrow Y$ have a bounded left inverse $L : Y \rightarrow E$. For every continuous coordinate functional $\ell \in E^*$, the composition

$$\alpha_{\ell} := \ell \circ L \in Y^*$$

satisfies

$$\alpha_{\ell}(DF(u_0)h) = \ell(h).$$

Thus every recovered infinitesimal latent coordinate has a dual quote functional that extracts it from the first-order option-manifold response.

Proof. The identity is the definition of a left inverse. If $LDF(u_0) = I_E$, then for every $h \in E$,

$$\alpha_\ell(DF(u_0)h) = \ell(LDF(u_0)h) = \ell(h).$$

Continuity follows from $\ell \in E^*$ and boundedness of L . $\square$

Corollary 7.8 (One index identifies one synchronization aggregate). *If the active unordered edge set satisfies $|\mathcal{E}| > 1$ and the only joint coordinate is one index or basket u , then $W_\mathcal{E}$ has one row. It cannot identify the full edge vector $V_\mathcal{E}(\tau)$. It identifies only the aggregate*

$$2 \sum_{(i,j) \in \mathcal{E}} w_i^u w_j^u V_{ij,t}^\mathbb{Q}(\tau).$$

Full cross-route recovery therefore requires additional carrier baskets, sector or index panels, a smaller active edge set, or source-family restrictions that reduce the effective dimension.

Proof. With one basket, $W_\mathcal{E}$ has a single row, so its rank is at most one. When $|\mathcal{E}| > 1$, the map from the active edge vector to public basket variance has a nontrivial kernel. The only identified cross term is the displayed row action of $W_\mathcal{E}$ on $V_\mathcal{E}(\tau)$. Separating the remaining edge coordinates requires more independent rows or a source restriction that lowers the quotient dimension. $\square$

Corollary 7.9 (Generic basket designs). *Suppose $r \geq |\mathcal{E}|$ basket weight vectors are drawn from a distribution with density on an open subset of $\mathbb{R}^N$, and suppose the design class contains one choice of weights for which $W_\mathcal{E}$ has rank $|\mathcal{E}|$. Then*

$$\text{rank } W_\mathcal{E} = |\mathcal{E}|$$

with probability one.

Proof. Some $|\mathcal{E}| \times |\mathcal{E}|$ minor determinant is then a nonzero polynomial in the basket weights. The zero set of a nonzero polynomial has Lebesgue measure zero, and a distribution with density assigns that set probability zero. $\square$

Definition 7.10 (Variance-strip covariance matrix). Fix a maturity window τ , and let $X_i(\tau)$ be the centered integrated return, variance, or volatility factor whose basket variance is observed. The variance-strip covariance matrix is

$$C(\tau) = (C_{ij}(\tau))_{i,j=1}^N, \quad C_{ij}(\tau) = \text{Cov}^\mathbb{Q}(X_i(\tau), X_j(\tau) \mid \mathcal{H}_t).$$

For a basket weight $w \in \mathbb{R}^N$, the model-free quadratic basket response is

$$V_w(\tau) = w^\top C(\tau) w.$$

Theorem 7.11 (Quadratic basket manifold and covariance recovery). *For a fixed τ , knowledge of $V_w(\tau)$ for all $w \in \mathbb{R}^N$ is equivalent to knowledge of the full symmetric covariance matrix $C(\tau)$. In particular,*

$$C_{ij}(\tau) = \frac{1}{2} \left(V_{e_i+e_j}(\tau) - V_{e_i}(\tau) - V_{e_j}(\tau) \right).$$

A finite family of basket weights $w_1, \dots, w_m$ identifies $C(\tau)$ on a target subspace $\mathcal{V} \subseteq \mathbb{S}^N$ if and only if the linear functionals

$$C \mapsto \text{Tr}(C w_a w_a^\top), \quad a = 1, \dots, m,$$

separate points of $\mathcal{V}$.

Proof. The map $C \mapsto w^\top C w = \text{Tr}(C w w^\top)$ is linear in the symmetric matrix C . Polarization gives the displayed identity. The finite-weight statement is exactly injectivity of this linear observation map on $\mathcal{V}$. $\square$

Theorem 7.12 (Elliptope feasibility and Perron uncertainty). *Suppose the single-name variance levels $C_{ii}(\tau) > 0$ are known, and set*

$$R_{ij}(\tau) = \frac{C_{ij}(\tau)}{\sqrt{C_{ii}(\tau) C_{jj}(\tau)}}.$$

A candidate cross-covariance matrix is statically feasible only if

$$R(\tau) \succeq 0, \quad R_{ii}(\tau) = 1.$$

If only a finite basket family is observed, the model-free feasible set is

$$\mathcal{C}_\tau := \{C \succeq 0 : \text{diag}(C) = v, \text{Tr}(C w_a w_a^\top) = V_{w_a}, a = 1, \dots, m\}.$$

For any screened Perron matrix $A(C)$ constructed from C , the model-free spectral-radius uncertainty interval is

$$[\rho_\tau^-, \rho_\tau^+] := \left[\inf_{C \in \mathcal{C}_\tau} \rho(A(C)), \sup_{C \in \mathcal{C}_\tau} \rho(A(C)) \right].$$

A Perron phase classification from the finite basket panel is model-free certified only when this interval lies entirely on one side of the phase threshold.

Proof. Positive semidefiniteness and unit diagonal are necessary for any correlation matrix and sufficient for a centered Gaussian vector with that correlation matrix. The finite basket constraints are affine in C , so the feasible set is an affine slice of the positive semidefinite cone. If the panel is feasible, the constraints $\text{diag}(C) = v$ with positive variance levels make this slice closed and bounded, hence compact. The map $C \mapsto \rho(A(C))$ is continuous for the screened Perron matrices considered here, so the infimum and supremum are attained on $\mathcal{C}_\tau$. If the displayed spectral interval crosses a phase threshold, two statically feasible covariance completions imply different phase labels. $\square$

Definition 7.13 (Direction-resolving response block). Let $G_{i \leftarrow j}$ denote the ordered Volterra route from source j to target i , and let $\vec{\mathcal{E}}$ be an active ordered edge set. A public observation block is direction-resolving on $\vec{\mathcal{E}}$ if the derivative of its response map with respect to the ordered route coordinates,

$$J_{\text{dir}} : (\delta G_{i \leftarrow j})_{(i,j) \in \vec{\mathcal{E}}} \longmapsto \delta \mathcal{M}_t^{\text{opt}},$$

is injective after quotienting by driver-basis gauge.

Theorem 7.14 (Symmetric covariance does not identify arbitrary direction). Suppose the admitted public response depends on ordered routes only through symmetric combinations

$$S_{ij} = \Psi_{ij}(G_{i \leftarrow j}, G_{j \leftarrow i}) = \Psi_{ij}(G_{j \leftarrow i}, G_{i \leftarrow j})$$

after the remaining source coordinates are fixed. Then any tangent direction satisfying

$$D\Psi_{ij}[\delta G_{i \leftarrow j}, \delta G_{j \leftarrow i}] = 0$$

for every active unordered pair lies in the kernel of the public observability operator. In particular, if locally $S_{ij} = G_{i \leftarrow j} + G_{j \leftarrow i}$, the anti-symmetric perturbation

$$\delta G_{i \leftarrow j} = H, \quad \delta G_{j \leftarrow i} = -H$$

is invisible to a symmetric covariance manifold.

Proof. If the public response factors through the symmetric map $S = (S_{ij})$, then its derivative factors as $D\bar{F} \circ DS$. Any tangent vector in $\text{Ker } DS$ is therefore in the kernel of the public response derivative. In the additive example, $DS[H, -H] = H - H = 0$, so the anti-symmetric perturbation lies in that kernel. $\square$

Theorem 7.15 (Sufficient orientation separation). Suppose that, after gauge normalization and after imposing all symmetric covariance constraints, each ordered route response on the active set admits the analytic representation

$$R_{i \leftarrow j}(\tau) = \sum_{q=1}^{d_{ij}} \gamma_{ijq} f_{ijq}(\tau), \quad \tau \in I,$$

where the family $\{f_{ijq} : (i, j, q) \in \vec{\mathcal{E}}\}$ is linearly independent on the admitted maturity interval I . Then an open maturity interval identifies the ordered route coefficients γ_{ijq} . A generic finite analytic maturity panel with sufficiently many points identifies the same coefficients.

Proof. The derivative of the response with respect to γ_{ijq} is the analytic function f_{ijq} . If a quotient tangent vector is annihilated on the open interval I , the corresponding linear combination of the functions f_{ijq} vanishes on I . Linear independence forces all tangent coefficients to be zero, so the continuum derivative is injective. The finite-panel conclusion follows from Corollary 8.17. $\square$

Definition 7.16 (Conditional basket cumulant tensor). For a return vector $\Delta_\tau X := X_{t+\tau} - X_t$, suppose the conditional moment generating function is finite near the origin. Define

$$\kappa_t(z, \tau) := \log \mathbb{E}^{\mathbb{Q}}[e^{z^\top \Delta_\tau X} \mid \mathcal{H}_t].$$

The order- n conditional cumulant tensor is

$$K_t^{(n)}(\tau)_{i_1 \dots i_n} := \partial_{z_{i_1}} \cdots \partial_{z_{i_n}} \kappa_t(0, \tau).$$

A full signed basket smile for $z^\top \Delta_\tau X$ determines $\kappa_t(z, \tau)$ near zero and hence all finite cumulant tensors whose moments exist.

Proposition 7.17 (Skew-tensor orientation test). *In a Volterra leverage source family, suppose an ordered edge $j \rightarrow i$ contributes to the short-maturity third cumulant through*

$$K_t^{(3)}(\tau)_{iij} = A_{i \leftarrow j} \tau^{\alpha_{i \leftarrow j}} + o(\tau^{\alpha_{i \leftarrow j}}), \quad A_{i \leftarrow j} \neq 0,$$

while the reverse edge contributes to

$$K_t^{(3)}(\tau)_{jji} = A_{j \leftarrow i} \tau^{\alpha_{j \leftarrow i}} + o(\tau^{\alpha_{j \leftarrow i}}).$$

If the pairs $(\alpha_{i \leftarrow j}, A_{i \leftarrow j})$ and $(\alpha_{j \leftarrow i}, A_{j \leftarrow i})$ are separated by the source-family gauge, then the observed short-maturity skew tensor identifies the orientation of the edge. If the third cumulant vanishes by symmetry, the same statement holds with the first nonzero mixed cumulant tensor.

Proof. The short-maturity power and leading coefficient are recovered from the cumulant response by the same leading-jet argument as Corollary 4.5. Separation of the two ordered leading pairs assigns the observed mixed cumulant to one orientation rather than the other. If the third cumulant is zero, the first nonzero mixed cumulant supplies the leading jet. $\square$

Theorem 7.18 (Minimal option-manifold hierarchy). *In a gauge-normalized source family with block tangent decomposition*

$$\bar{E}^{\text{str}} = E_{\mathfrak{S}} \oplus E_G \oplus E_{\text{syn}},$$

the following hierarchy holds locally.

- (1) *Single-name vanilla term structures can identify marginal maturity-scaling directions in $E_{\mathfrak{S}}$, but leave $E_G \oplus E_{\text{syn}}$ unobserved unless cross-kernel coordinates are fixed by the source family.*
- (2) *Adding nondegenerate index or basket option responses creates a synchronization observation block for E_{syn} .*
- (3) *Adding a nondegenerate carrier span with Gram lower bound $\kappa_{\text{car}} > 0$ separates carrier-replicated single-name variation from the synchronized index component.*
- (4) *An open maturity interval, or a finite analytic maturity panel with full rank, upgrades synchronization curves into recovery of rough maturity and Volterra source coordinates.*
- (5) *Residual-selector coordinates lie outside the purely public source tangent space. They remain unobserved unless residual price instruments are included or a selector restriction is imposed.*

Thus full canonical recovery requires a joint index/single-name/carrier/maturity manifold beyond calibrated marginal surfaces.

Proof. Each statement follows by restricting the observability operator $\bar{\mathcal{O}}_t^{\mathbb{Q}}$ to the indicated public subsystem. Marginal vanilla payoffs differentiate only marginal law and maturity-scaling directions, so the joint blocks remain in the kernel. Index and basket responses contain cross terms and therefore generate a synchronization block. Carrier Gram nondegeneracy makes the carrier projection invertible on the admitted carrier coordinates, separating replicated single-name variation from common synchronization. The finite-panel assertion is the analytic rank argument later formalized in Theorem 8.16. Finally, residual selectors change prices only along

public-equivalent residual directions; absent residual instruments or a selector restriction, those directions remain outside the public observation map. $\square$

Definition 7.19 (Finite-quote observation error). A noisy finite quote panel is

$$\hat{y}_t = y_t + \varepsilon_t, \quad y_t = \mathcal{M}_t^{\text{opt}},$$

where y_t lies in a recovery neighborhood and $\varepsilon_t \in \mathcal{Y}_t$ is a quote, smoothing, or finite-panel perturbation measured in the quote-space norm.

Theorem 7.20 (Finite-quote stability and observability breakdown). *Suppose y_t and $\hat{y}_t$ remain in the same local recovery neighborhood and $\hat{y}_t \in W_t \cap F_t(V_t)$, so that the local inverse is defined. Let*

$$\hat{\Xi}_t^{\mathbb{Q}} = \mathcal{I}_t(\hat{y}_t), \quad \Xi_t^{\mathbb{Q}} = \mathcal{I}_t(y_t).$$

Then

$$\text{dist}_{\bar{\mathcal{X}}_t}(\hat{\Xi}_t^{\mathbb{Q}}, \Xi_t^{\mathbb{Q}}) \leq 2 \|L_t\| \|\varepsilon_t\|_{\mathcal{Y}_t}.$$

In finite-dimensional panels, if

$$\beta_t := \sigma_{\min}(\bar{\mathcal{O}}_t^{\mathbb{Q}}) > 0,$$

then $\|L_t\| \leq \beta_t^{-1}$. Hence recovery becomes ill-conditioned when $\beta_t \downarrow 0$. In the Volterra–Perron chart this can occur through collapse of the Perron gap or reduced-resolvent conditioning, carrier Gram lower bound, maturity-response rank, directed-route rank, synchronization rank, or, in an augmented residual-instrument manifold, residual-selector rank.

Proof. The first estimate is the local inverse estimate from Theorem 8.4 applied to y_t and $y_t + \varepsilon_t$. In a finite-dimensional panel, the Moore–Penrose left inverse of the quotient observability matrix has operator norm $\sigma_{\min}(\bar{\mathcal{O}}_t^{\mathbb{Q}})^{-1}$ on the observed range. The block decomposition then identifies the ways in which this singular value can vanish. Rank loss in a diagonal block, divergence of the Perron reduced-resolvent constant, collapse of the carrier Gram bound, or loss of a residual-selector rank block removes a bounded left inverse and forces the local recovery constant to diverge. $\square$

Remark 7.21. If a noisy quote vector does not lie on the model image, the inverse $\mathcal{I}_t$ is not defined. In that case the regularized finite-panel estimator of Definition 8.19 is the appropriate reconstruction object.

Definition 7.22 (Bid–ask quote set). For a finite panel, let $\underline{y}_t, \bar{y}_t \in \mathbb{R}^M$ be lower and upper executable quote vectors. Let $\mathcal{A}$ denote the finite set of static-arbitrage linear inequalities used to filter the panel. The arbitrage-filtered quote set is

$$\mathcal{Q}_t := \{y \in \mathbb{R}^M : \underline{y}_t \leq y \leq \bar{y}_t, \mathcal{A}y \geq 0\}.$$

For a finite-panel response map $F_\tau : U \rightarrow \mathbb{R}^M$, the set of model states compatible with the quote set is

$$\mathcal{I}_t(\mathcal{Q}_t) := \{u \in U : F_\tau(u) \in \mathcal{Q}_t\}.$$

Theorem 7.23 (Set-valued recovery under bid–ask uncertainty). *Suppose F_τ has a Lipschitz inverse on the chosen recovery neighborhood with constant C_τ . Then*

$$\text{diam}(\mathcal{I}_t(\mathcal{Q}_t)) \leq C_\tau \text{diam}(\mathcal{Q}_t \cap F_\tau(U)).$$

If the whole bid–ask box has quote-space diameter at most δ_t , then

$$\text{diam}(\mathcal{I}_t(\mathcal{Q}_t)) \leq C_\tau \delta_t.$$

For any selected claim price $\Pi(X; u)$ that is Lipschitz in the recovered state with constant L_X ,

$$\text{diam}\{\Pi(X; u) : u \in \mathcal{I}_t(\mathcal{Q}_t)\} \leq L_X C_\tau \delta_t,$$

plus the residual-envelope width if residual selectors are allowed to vary.

Proof. For any $u, v \in \mathcal{I}_t(\mathcal{Q}_t)$, both $F_\tau(u)$ and $F_\tau(v)$ belong to $\mathcal{Q}_t \cap F_\tau(U)$. The inverse estimate gives

$$\|u - v\| \leq C_\tau \|F_\tau(u) - F_\tau(v)\|.$$

Taking the supremum over compatible states gives the diameter bound. If the whole filtered bid–ask box has diameter at most δ_t , the same estimate gives the second bound. Applying the Lipschitz property of $\Pi(X; u)$ to the recovered-state set gives the price-diameter bound; any free selector coordinate adds its residual-envelope width separately. $\square$

8. LATENT-STATE RECOVERY FROM RICH OPTION MANIFOLDS

Restricted source families can restore recovery despite global non-identification. Public compression alone is insufficient; observability supplies the additional condition. A resolving option manifold contains enough cross-sectional and maturity variation to recover the source quotient inside the specified family.

Assumption 8.1 (Gauge-normalized source family). There is a Banach manifold $\mathfrak{X}_t^{\text{str}}$ of admissible Volterra–Perron states and a gauge-normalized local chart

$$\psi : U \subset E \rightarrow \mathfrak{X}_t^{\text{str}}, \quad E = E_{\mathfrak{H}} \oplus E_G \oplus E_\rho \oplus E_e \oplus E_P.$$

The local gauge-slice condition in Assumption 3.12 holds. Thus gauge-equivalent infinitesimal directions form a closed subspace $E_g \subset E$, and the canonical tangent space is identified with the slice $E_c \simeq E/E_g =: \bar{E}$. The chart represents the coordinates

$$(\mathfrak{H}^\mathbb{Q}, G^\mathbb{Q}, \rho^\mathbb{Q}, e^\mathbb{Q}, \ell^\mathbb{Q}, P^\mathbb{Q})$$

using the normalizations of Definition 3.14.

Assumption 8.2 (Observable response rank). At the reference state $\Xi_0 = \psi(u_0)$, the observable response map

$$F := \mathcal{R}_t \circ \psi : U \rightarrow \mathcal{Y}_t$$

is continuously Fréchet differentiable. Its derivative annihilates gauge directions and induces

$$\overline{DF}(u_0) : \bar{E} \rightarrow \text{Ran } DF(u_0) \subseteq \mathcal{Y}_t.$$

This induced derivative has a bounded left inverse $L_0 : \text{Ran } DF(u_0) \rightarrow \bar{E}$:

$$L_0 \overline{DF}(u_0) = \text{Id}_{\bar{E}}.$$

Equivalently, with $B := \|L_0\|$,

$$\|h\|_{\bar{E}} \leq B \left\| \overline{DF}(u_0) h \right\|_{\mathcal{Y}_t}, \quad h \in \bar{E}.$$

Assumption 8.3 (Rank and conditioning). The reference state satisfies the following lower-bound conditions.

- (1) The screened synchronization operator has a simple Perron eigenvalue with spectral gap $\gamma_P > 0$ and finite reduced-resolvent constant $\kappa_P := \left\| (A_t^\mathbb{Q} - \rho_t^\mathbb{Q} I)_{\text{red}}^{-1} \right\|$ in the chosen Perron chart.
- (2) The admitted carrier coordinates have Gram matrix Γ_{car} satisfying $\langle x, \Gamma_{\text{car}} x \rangle \geq \kappa_{\text{car}} \|x\|^2$ for some $\kappa_{\text{car}} > 0$.
- (3) The rough maturity-response block has full rank either on an open maturity interval or on a finite analytic maturity panel.
- (4) The directed Volterra source-family block and the index/single-name synchronization block have full rank after quotienting by gauge.
- (5) Any residual selector used for pricing is either fixed by a post-recovery rule or observed only through an augmented residual-instrument block, not through the pure public option response.

Theorem 8.4 (Local identification of the canonical Volterra–Perron source quotient). *Under Assumptions 8.1–8.3, there are neighborhoods $V \subset U$ of u_0 and $W \subset \mathcal{Y}_t$ of $F(u_0)$ such that the public option manifold identifies the canonical source state locally:*

$$F(u) = F(v), \quad u, v \in V \quad \implies \quad \psi(u) \sim_g \psi(v).$$

Equivalently, the quotient map

$$\bar{F} : V/E_g \rightarrow W \cap F(V)$$

has a continuous local inverse on its image. After shrinking V , the inverse obeys a local Lipschitz estimate

$$\|[u] - [v]\|_{\bar{E}} \leq 2 \|L_0\| \|F(u) - F(v)\|_{\mathcal{Y}_t}, \quad u, v \in V.$$

Proof. Because $DF(u_0)$ annihilates gauge directions and the response is constant on local gauge orbits, F descends locally to a differentiable map $\bar{F}$ on the quotient chart. Quotient representatives are again denoted by u, v . Set

$$J := D\bar{F}(u_0).$$

By the bounded-below form of Assumption 8.2,

$$\|h\|_{\bar{E}} \leq B \|Jh\|_{\mathcal{Y}_t}, \quad h \in \bar{E}.$$

Continuity of $D\bar{F}$ gives a convex neighborhood V on which

$$\left\| D\bar{F}(z) - J \right\|_{\mathcal{L}(\bar{E}, \mathcal{Y}_t)} \leq \frac{1}{2B}.$$

For $u, v \in V$, the fundamental theorem of calculus in the Banach chart gives

$$\bar{F}(u) - \bar{F}(v) = J(u - v) + R(u, v), \quad R(u, v) := \int_0^1 (D\bar{F}(v + s(u - v)) - D\bar{F}(u_0))(u - v) ds.$$

The neighborhood choice implies

$$\|R(u, v)\|_{\mathcal{Y}_t} \leq \frac{1}{2B} \|u - v\|_{\bar{E}}.$$

Therefore

$$\begin{aligned} \|u - v\|_{\bar{E}} &\leq B \|J(u - v)\|_{\mathcal{Y}_t} \\ &\leq B \left\| \bar{F}(u) - \bar{F}(v) \right\|_{\mathcal{Y}_t} + B \|R(u, v)\|_{\mathcal{Y}_t} \\ &\leq B \left\| \bar{F}(u) - \bar{F}(v) \right\|_{\mathcal{Y}_t} + \frac{1}{2} \|u - v\|_{\bar{E}}. \end{aligned}$$

This gives the displayed Lipschitz estimate and hence local injectivity on the quotient. If $F(u) = F(v)$, then $[u] = [v]$, so $\psi(u) \sim_g \psi(v)$. The same estimate gives continuity of the inverse on $F(V)$. $\square$

Remark 8.5. The theorem is conditional and compatible with the non-identification theorem. Recovery becomes possible only after restricting the latent state to a source family whose option responses are locally complete. Public-annihilating residual selectors remain unidentified from the same public quotes.

Theorem 8.6 (Stability of recovered states). *Suppose Assumptions 8.1–8.3 hold and let*

$$B := \|L_0\| < \infty.$$

Then, after shrinking the recovery neighborhood if necessary, there are constants $C > 0$ and $\delta > 0$ such that whenever two option manifolds $y_1, y_2 \in W \cap F(V)$ satisfy $\|y_1 - y_2\|_{\mathcal{Y}_t} < \delta$, their recovered canonical states satisfy

$$\text{dist}_{\bar{E}}(\bar{F}^{-1}(y_1), \bar{F}^{-1}(y_2)) \leq C \|y_1 - y_2\|_{\mathcal{Y}_t}.$$

One may take $C = 2B$ on a sufficiently small neighborhood. In the Volterra–Perron chart, B is finite only while the Perron reduced-resolvent constant κ_P remains bounded, the carrier Gram lower bound κ_{car} remains bounded away from zero, and the diagonal response-block singular values remain bounded away from zero.

Proof. The first estimate is the Lipschitz inverse bound proved in Theorem 8.4. The Perron contribution is controlled by the reduced resolvent of the simple eigenvalue; it can blow up as the spectral gap closes or as non-normal conditioning degenerates. The dependence on κ_{car} follows because the carrier coordinates are recovered by inverting the carrier Gram matrix. The remaining diagonal response blocks enter through the norm of the block left inverse. $\square$

Definition 8.7 (Observation modulus). Let $F : U \subset E \rightarrow Y$ be a quotient response map on a recovery chart. For $r > 0$, define

$$\omega_F(r) := \inf\{\|F(u) - F(v)\|_Y : u, v \in U, \|u - v\|_E \geq r\}.$$

For a finite linear panel $F(u) = F(u_0) + J(u - u_0)$ with full column-rank derivative, $\omega_F(r) = r \sigma_{\min}(J)$ locally.

Theorem 8.8 (No reconstruction beats the inverse modulus). *Suppose only a perturbed quote vector $\hat{y}$ is observed, with deterministic error bound*

$$\|\hat{y} - F(u)\|_Y \leq \varepsilon.$$

Let $\hat{u}(\hat{y})$ be any reconstruction rule. If there exist $u, v \in U$ such that

$$\|u - v\|_E \geq r, \quad \|F(u) - F(v)\|_Y \leq 2\varepsilon,$$

then, for at least one admissible observation $\hat{y}$,

$$\max\{\|\hat{u}(\hat{y}) - u\|_E, \|\hat{u}(\hat{y}) - v\|_E\} \geq \frac{r}{2}.$$

Consequently, if $\omega_F(r) \leq 2\varepsilon$, the worst-case recovery error at quote-error level ε is at least $r/2$.

Proof. In a normed linear quote space take $\hat{y} = (F(u) + F(v))/2$. The bound $\|F(u) - F(v)\|_Y \leq 2\varepsilon$ makes $\hat{y}$ admissible for both latent states. The same observed quote vector is therefore compatible with u and with v . If one output $\hat{u}(\hat{y})$ were within $r/2$ of both states, the triangle inequality would give $\|u - v\|_E < r$, contradicting the separation condition. Hence at least one of the two errors is at least $r/2$. $\square$

Corollary 8.9 (Sharp local recovery scale). *For a finite panel with least quotient singular value $\beta > 0$, deterministic quote uncertainty ε implies an unavoidable local recovery scale of order*

$$\frac{\varepsilon}{\beta}.$$

Thus, as $\beta \downarrow 0$, recovery instability is an information-theoretic limit for every procedure using the same quote data.

Proof. On the quotient tangent space, the least singular value β gives the smallest first-order quote displacement per unit latent displacement. Taking a unit singular direction associated with β , two nearby latent states separated by order ε/β produce quote vectors separated by order ε , up to higher-order terms. The local minimax lower bound in Theorem 8.8 then gives the same scale, while the inverse estimate from the bounded left inverse gives the matching upper order. $\square$

Corollary 8.10 (Recovery and residual freedom are complementary). *In a locally complete source family, the public option manifold recovers the canonical Volterra–Perron state. Outside such a family, public-equivalent residual density families remain available and the residual price interval of [12] is the correct pricing object.*

Proof. The first statement is Theorem 8.4. The second is Theorem 6.7. They apply to different information sets, namely a restricted source family with full response rank and an unrestricted residual kernel class. $\square$

8.1. Block recovery and finite maturity panels. The quotient inverse theorem separates recovery into response blocks. In the option manifold, these blocks have financial content. Marginal single-name term structures carry rough maturity geometry; index-versus-carrier spreads carry synchronization. A residual charge becomes a recoverable coordinate only in an augmented manifold containing residual instruments.

Definition 8.11 (Observable block response). The observable response map is block-adapted if, after a local linear recombination of public quote coordinates, its response space decomposes into

$$\mathcal{Y}_t = \mathcal{Y}_H \oplus \mathcal{Y}_G \oplus \mathcal{Y}_{\text{syn}}.$$

The blocks are interpreted as maturity-scaling responses, directed cross-kernel responses, and index-versus-carrier synchronization responses. An augmented residual-instrument manifold may append an additional block $\mathcal{Y}_\theta$.

Definition 8.12 (Block source chart). A source chart is block triangular if its quotient tangent space decomposes as

$$\bar{E} = E_{\mathcal{J}} \oplus E_G \oplus E_{\text{syn}}, \quad E_{\text{syn}} := E_\rho \oplus E_e \oplus E_P,$$

the observable response is block-adapted, so that the derivative of the pricing map has the block form

$$\overline{DF}(u_0) = \begin{pmatrix} J_H & 0 & 0 \\ * & J_G & 0 \\ * & * & J_{\text{syn}} \end{pmatrix}.$$

The blocks correspond respectively to rough maturity scaling, directed Volterra/cross-kernel geometry, and Perron synchronization.

Theorem 8.13 (Block-rank observability criterion). *Suppose the source chart is block triangular and each diagonal block J_H, J_G, J_{syn} has a bounded left inverse on its range. Then $\overline{DF}(u_0)$ has a bounded left inverse on its range. Consequently the conditional recovery theorem applies to the canonical state. If an augmented residual-instrument block is present and its diagonal block J_θ also has a bounded left inverse, the same triangular argument recovers the selected residual coordinate as part of the augmented state.*

Proof. Write a tangent vector as $h = (h_H, h_G, h_{\text{syn}})$. Given the response $y = \overline{DF}(u_0)h$, recover h_H from the first block by applying the left inverse of J_H . Subtract the known contribution of h_H from the second block and recover h_G by the left inverse of J_G . Continue recursively through the synchronization block. The finite number of bounded operations gives a bounded left inverse for the full triangular derivative. Theorem 8.4 then gives local recovery. $\square$

Theorem 8.14 (Triangular block inverse bound). *Let*

$$J = \begin{pmatrix} J_1 & 0 & \cdots & 0 \\ C_{21} & J_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ C_{m1} & C_{m2} & \cdots & J_m \end{pmatrix}$$

act between finite products of normed block spaces. Suppose each diagonal block has a bounded left inverse L_i with $\|L_i\| \leq b_i$, and set $c_{ij} := \|C_{ij}\|$ for $j < i$. Define

$$B_1 = b_1, \quad B_i = b_i \left(1 + \sum_{j < i} c_{ij} B_j \right), \quad i = 2, \dots, m.$$

Then J has a bounded left inverse on its range, and one may choose it so that

$$\|J_{\text{left}}^{-}\| \leq \left(\sum_{i=1}^m B_i^2 \right)^{1/2}$$

under product norms. In the Volterra–Perron chart, the b_i ’s are controlled by the inverse rough-maturity rank, directed-route rank, Perron reduced-resolvent conditioning, carrier Gram lower bound, synchronization rank, and any augmented residual-selector rank.

Proof. Given $y = Jh$, solve recursively. First $h_1 = L_1 y_1$, so $\|h_1\| \leq b_1 \|y\|$. For $i \geq 2$,

$$h_i = L_i \left(y_i - \sum_{j < i} C_{ij} h_j \right).$$

Thus

$$\|h_i\| \leq b_i \left(\|y\| + \sum_{j < i} c_{ij} \|h_j\| \right).$$

The displayed recursion gives $\|h_i\| \leq B_i \|y\|$. Summing over blocks gives the claimed operator norm bound. $\square$

Definition 8.15 (Finite maturity observation operator). Let $I \subset (0, \infty)$ be an open maturity interval and let

$$R_\Theta : I \rightarrow \mathbb{R}^m$$

be a vector of option-response functions generated by a state Θ . For a finite maturity set $\tau = (\tau_1, \dots, \tau_k) \subset I$, define

$$\mathcal{O}_\tau(\Theta) := (R_\Theta(\tau_1), \dots, R_\Theta(\tau_k)) \in \mathbb{R}^{mk}.$$

Theorem 8.16 (Finite panels recover analytic source charts). *Assume the source chart is finite dimensional after quotienting by gauge, and for each state the response function $\tau \mapsto R_\Theta(\tau)$ is real analytic on I . If the continuum response map*

$$\Theta \mapsto \{R_\Theta(\tau) : \tau \in I\}$$

has full local rank at Θ_0 , then there exists a finite maturity set $\tau \subset I$ such that the finite observation operator $\mathcal{O}_\tau$ has full local rank at Θ_0 . Hence a finite index/single-name maturity panel locally recovers the canonical source state.

Proof. Let d be the dimension of the quotient source chart. Full rank of the continuum response means that the d derivative functions

$$\tau \mapsto DR_{\Theta_0}(\tau)[v_j], \quad j = 1, \dots, d,$$

are linearly independent for some basis $v_1, \dots, v_d$ of the quotient tangent space. Let $(q_n)_{n \geq 1}$ be a countable dense subset of I . If every finite maturity evaluation were rank deficient, then for each n there would be a vector $a^{(n)} \in \mathbb{R}^d$, $\|a^{(n)}\| = 1$, such that

$$\sum_{j=1}^d a_j^{(n)} DR_{\Theta_0}(q_k)[v_j] = 0, \quad k = 1, \dots, n.$$

By compactness of the unit sphere, pass to a subsequence $a^{(n_\ell)} \rightarrow a$ with $\|a\| = 1$. For each fixed k , the preceding display holds for all sufficiently large ℓ , hence

$$\sum_{j=1}^d a_j DR_{\Theta_0}(q_k)[v_j] = 0.$$

The analytic function

$$\tau \mapsto \sum_{j=1}^d a_j DR_{\Theta_0}(\tau)[v_j]$$

therefore vanishes on a dense subset of I , and hence identically. This contradicts linear independence. Therefore some finite maturity set has full rank. The finite-dimensional inverse function theorem gives local recovery from that finite panel. $\square$

Corollary 8.17 (Generic finite analytic panels). *Under the assumptions of Theorem 8.16, there exist k and a $d \times d$ minor of the finite-panel Jacobian whose determinant is a nonzero real-analytic function of $(\tau_1, \dots, \tau_k) \in I^k$. Hence the set of k -point maturity panels with deficient local rank is contained in the zero set of a nonzero real-analytic function and has Lebesgue measure zero in I^k .*

Proof. The proof of Theorem 8.16 gives at least one finite panel whose Jacobian has rank d . Therefore some $d \times d$ minor is nonzero at that panel. The entries of the finite-panel Jacobian are analytic in the maturity coordinates, so this minor is a nonzero analytic function. The zero set of a nonzero real-analytic function has Lebesgue measure zero. $\square$

Corollary 8.18 (Finite-panel observability determinant). *For a fixed finite panel τ , let J_τ be the quotient finite-panel Jacobian at the reference state. If the quotient chart has dimension d and J_τ has at least d quote rows, define*

$$\mathfrak{D}_t(\tau) := \det(J_\tau^* J_\tau).$$

Then $\mathfrak{D}_t(\tau) > 0$ is equivalent to local finite-panel observability. When this holds,

$$\|J_\tau^-\| \leq \sigma_{\min}(J_\tau)^{-1} = \mathfrak{D}_t(\tau)^{-1/2} \prod_{j=1}^{d-1} \sigma_j(J_\tau),$$

where $\sigma_1(J_\tau) \geq \dots \geq \sigma_d(J_\tau) > 0$ are the singular values. Loss of recovery is therefore visible as the finite-panel observability determinant approaching zero.

Proof. The quotient finite-panel map is locally observable exactly when its Jacobian has rank d . This is equivalent to $J_\tau^* J_\tau$ being positive definite, hence to $\mathfrak{D}_t(\tau) > 0$. If $\sigma_1 \geq \dots \geq \sigma_d > 0$ are the singular values, then

$$\mathfrak{D}_t(\tau) = \prod_{j=1}^d \sigma_j(J_\tau)^2.$$

The Moore–Penrose left inverse has norm $\sigma_d(J_\tau)^{-1}$, which gives the displayed identity after solving the determinant formula for σ_d^{-1} . Thus determinant collapse is exactly singular-value collapse in the quotient chart. $\square$

Definition 8.19 (Regularized finite-panel reconstruction). Let $F_\tau : \bar{U} \subset \bar{E} \rightarrow \mathbb{R}^M$ be a finite-panel observation map on the gauge quotient, and let

$$\hat{y}_\tau = F_\tau(u_0) + \varepsilon_\tau$$

be a finite quote vector with smoothing or observation error ε_τ . For a compact recovery neighborhood $K \subset \bar{U}$, a reference point $u_{\text{pr}} \in K$, and a regularization level $\alpha \geq 0$, define

$$\hat{u}_{\alpha,\tau} \in \arg \min_{u \in K} \left\{ \|F_\tau(u) - \hat{y}_\tau\|_{\mathbb{R}^M}^2 + \alpha \|u - u_{\text{pr}}\|_{\bar{E}}^2 \right\}.$$

The recovered risk-neutral state is $\hat{\Xi}_{\alpha,\tau}^{\mathbb{Q}} = \psi(\hat{u}_{\alpha,\tau})$, with the gauge normalizations fixed by the chart.

Theorem 8.20 (Regularized finite-panel recovery). *Suppose F_τ is continuous on the compact recovery neighborhood K and has a local inverse on K with Lipschitz constant C_τ :*

$$\|u - v\|_{\bar{E}} \leq C_\tau \|F_\tau(u) - F_\tau(v)\|_{\mathbb{R}^M}, \quad u, v \in K.$$

Let $u_0 \in K$, let $\hat{u}_{\alpha,\tau}$ be any minimizer from Definition 8.19, and set

$$r_{\text{pr}} := \|u_0 - u_{\text{pr}}\|_{\bar{E}}.$$

Then, whenever the minimizer remains in the same recovery neighborhood,

$$\|\hat{u}_{\alpha,\tau} - u_0\|_{\bar{E}} \leq C_\tau \left(\|\varepsilon_\tau\|_{\mathbb{R}^M} + \sqrt{\|\varepsilon_\tau\|_{\mathbb{R}^M}^2 + \alpha r_{\text{pr}}^2} \right) \leq C_\tau (2 \|\varepsilon_\tau\|_{\mathbb{R}^M} + \sqrt{\alpha} r_{\text{pr}}).$$

If the finite-panel observability matrix has least quotient singular value $\beta_\tau > 0$, then C_τ may be chosen proportional to β_τ^{-1} on a sufficiently small neighborhood. Thus finite-panel recovery is stable only to the extent that quote error, smoothing error, regularization bias, and observability conditioning are controlled.

Proof. A minimizer exists because K is compact and the objective in Definition 8.19 is continuous on K . By optimality of $\hat{u}_{\alpha,\tau}$,

$$\|F_\tau(\hat{u}_{\alpha,\tau}) - \hat{y}_\tau\|^2 + \alpha \|\hat{u}_{\alpha,\tau} - u_{\text{pr}}\|^2 \leq \|F_\tau(u_0) - \hat{y}_\tau\|^2 + \alpha \|u_0 - u_{\text{pr}}\|^2 = \|\varepsilon_\tau\|^2 + \alpha r_{\text{pr}}^2.$$

Therefore

$$\|F_\tau(\hat{u}_{\alpha,\tau}) - \hat{y}_\tau\| \leq \sqrt{\|\varepsilon_\tau\|^2 + \alpha r_{\text{pr}}^2}.$$

The triangle inequality gives

$$\|F_\tau(\hat{u}_{\alpha,\tau}) - F_\tau(u_0)\| \leq \sqrt{\|\varepsilon_\tau\|^2 + \alpha r_{\text{pr}}^2} + \|\varepsilon_\tau\|.$$

Applying the local inverse estimate yields the first displayed bound. The second follows from $\sqrt{a^2 + b^2} \leq a + b$. The singular-value statement is the finite-dimensional inverse bound for the quotient observability matrix. $\square$

Corollary 8.21 (Sieve consistency of recovered source states). *Consider a sequence of finite panels τ_n , recovery neighborhoods K_n , and source charts whose approximation error to the target canonical state is a_n . If*

$$C_{\tau_n} (\|\varepsilon_{\tau_n}\| + \sqrt{\alpha_n} r_{\text{pr},n}) + a_n \longrightarrow 0,$$

then the reconstructed states converge to the target canonical risk-neutral state in quotient distance. In particular, when $C_{\tau_n} r_{\text{pr},n}$ is bounded, a convenient sufficient regularization schedule is

$$\alpha_n \downarrow 0, \quad \frac{\|\varepsilon_{\tau_n}\|^2}{\alpha_n} \longrightarrow 0, \quad C_{\tau_n} \|\varepsilon_{\tau_n}\| \longrightarrow 0, \quad a_n \longrightarrow 0.$$

The middle condition is the quotient-space analogue of the classical discrepancy scaling. The penalty vanishes asymptotically, but not so quickly that the reconstruction can spend the noise envelope as source motion.

Proof. Apply Theorem 8.20 inside each finite source chart and add the chart approximation error a_n by the triangle inequality. $\square$

Theorem 8.22 (Volterra–Perron option-manifold equivalence theorem). *Fix date t , an arbitrage-free public option manifold, an ambient latent state space $\mathfrak{X}_t$, and a gauge-normalized Volterra–Perron source family $\mathfrak{S}_t \subset \mathfrak{X}_t$. Let*

$$\bar{F}_t : \mathfrak{S}_t / \sim_g \longrightarrow \mathfrak{X}_t^{\text{pub}}$$

be the public response on the source gauge quotient. Around a reference source state, the following three statements are equivalent.

- (1) *The canonical $\mathbb{Q}$ -state is locally and stably recoverable from the public option manifold, up to the stated scale, sign, latent-basis, and Perron-normalization gauge conventions: the public response has a locally Lipschitz inverse on the model image of the recovery stratum.*
- (2) *The source quotient embeds locally and regularly into the public-observable quotient: $\bar{F}_t$ is a C^1 local embedding on the recovery stratum, its image is a local embedded submanifold of $\mathfrak{X}_t^{\text{pub}}$, and its inverse on the image is locally Lipschitz in the quote norm.*
- (3) *The quotient observability operator has no source kernel and has positive local modulus:*

$$\text{Ker } D\bar{F}_t = \{0\}, \quad \mathfrak{m}_t := \inf_{\|h\|_{\bar{E}}=1} \|D\bar{F}_t h\|_{\mathcal{Y}_t} > 0$$

on the recovery stratum.

In the block-triangular source chart, these equivalent conditions are verified by a full-rank diagonal-block certificate. The maturity/source-response block, directed covariance-route block, and index/carrier synchronization block each have a bounded left inverse on their quotient domains. Thus maturity scaling separates rough source geometry, carrier and basket responses separate active Volterra routes, and synchronization responses separate the simple Perron coordinate. When these equivalent conditions hold, the local inverse satisfies

$$\text{dist}_{\bar{E}}(\bar{F}_t^{-1}(y_1), \bar{F}_t^{-1}(y_2)) \leq C \mathfrak{m}_t^{-1} \|y_1 - y_2\|_{\mathcal{Y}_t}$$

for nearby public manifolds in the model image. For a finite panel with least quotient singular value $\beta_t > 0$, deterministic recovery error is at most order ε/β_t under quote error ε , and no reconstruction rule can improve this order uniformly.

The residual statement is separate from state recovery. For a latent loading class $\mathcal{L}_{t,T}$, selected residual prices collapse to singletons if and only if the admitted public payoff module is complete for that class:

$$\mathcal{L}_{t,T} \subseteq \mathcal{A}_{t,T}^{\text{pub}}.$$

Outside public completion, the remaining unidentified directions are exactly the public-annihilator residual pricing fibers generated by public-equivalent density families. Thus recovery of the Volterra–Perron state does not by itself imply residual-free pricing, tradeability, or hedgeability of every latent-stress claim. Dynamic admissibility further requires the static calendar convex-order restrictions and, in a recovered semimartingale chart, the martingale drift restrictions for admitted option coordinates.

Proof. The public option manifold identifies exactly $\mathfrak{X}_t^{\text{pub}}$ by Theorem 6.2. A recovery rule on the source family must therefore invert the induced map from $\mathfrak{S}_t/\sim_g$ into this public quotient. Local stable recovery is exactly local Lipschitz invertibility of that induced map on the model image, which is the embedded-stratum statement.

In a recovery chart, a bounded left inverse of the quotient derivative gives the local inverse estimate of Theorem 8.4. In finite dimension this is equivalent to a positive least quotient singular value; in the Banach chart it is the bounded-below hypothesis used there. Conversely, a C^1 local embedding with locally Lipschitz inverse has no quotient-source kernel and has positive local modulus on the recovery stratum. Singular analytic identifiability with zero derivative belongs to a seam and is treated separately by Theorem 8.27.

The block-rank condition is the triangular verification mechanism proved in Theorem 8.13. It constructs a bounded left inverse by solving successively for the maturity, directed-route, and synchronization coordinates. If a diagonal response block ceases to be bounded below and the lower blocks do not supply an independent response for that lost quotient direction, the public map has either a kernel or a vanishing inverse modulus along that direction. The finite-panel singular values record this obstruction in the block chart.

The stability estimate is Theorem 8.6; the finite-panel upper bound is Theorem 8.20; and the matching lower bound is Theorem 8.8. The residual clause is Theorem 6.11. Its proof identifies residual-free pricing with membership of the claim loading in the public payoff module and identifies non-uniqueness with the public-annihilator directions of Lemma 6.6. The dynamic admissibility clause combines Theorem 2.8 with the recovered-state martingale restriction proved later in Theorem 10.4. $\square$

Remark 8.23 (What the equivalence does not say). The theorem separates recoverability, residual-free pricing, and tradeable hedgeability:

$$\text{recoverability} \neq \text{residual-free pricing} \neq \text{tradeable hedgeability}.$$

The option manifold may recover the state while non-public claims still lie in residual pricing fibers. Conversely, a public claim can be priced or hedged without recovering the full Volterra–Perron state. This distinction is part of the mathematical-finance content of the theory rather than a limitation of the recovery theorem.

8.2. Failure modes. The equivalence theorem also classifies non-identification. Failures can occur at the public no-arbitrage layer, the source quotient, the carrier span, the Perron mode, the maturity dictionary, the residual payoff module, or the dynamic martingale layer.

Definition 8.24 (Failure certificates). For a finite option panel and a fixed source chart, the failure certificate vector is the collection

$$\mathfrak{C}_t := \left(\delta_{\text{arb}}, \delta_{\text{cx}}, \text{rank } W_{\mathcal{E}}, \lambda_{\min} \Gamma_{\text{car}}, \Delta_P, \kappa_P, \sigma_{\min} J_{\text{mat}}, \mathbf{m}_t, \dim \mathcal{N}_{t,T}^{\text{pub}}, \text{width } \mathcal{I}_t(X) \right).$$

Here δ_{arb} is distance to the static no-arbitrage cone, δ_{cx} is a convex-order violation across maturities, $W_{\mathcal{E}}$ is the basket design matrix, Γ_{car} is the carrier Gram matrix, Δ_P is the Perron

spectral gap, κ_P is the reduced-resolvent norm for the simple Perron eigenvalue, and J_{mat} is the maturity-response Jacobian,

$$\mathbf{m}_t := \inf_{\|h\|_{\bar{E}}=1} \left\| \overline{\mathcal{O}}_t^{\mathbb{Q}} h \right\|_{\mathcal{Y}_t}$$

is the quotient observability modulus, $\mathcal{N}_{t,T}^{\text{pub}}$ is the public annihilator, and $\mathcal{I}_t(X)$ is the residual price interval for X .

Theorem 8.25 (Failure-mode theorem). *Consider a finite-dimensional Volterra–Perron source chart satisfying the gauge-slice, smoothness, and no-arbitrage assumptions above. Stable dynamic selected pricing from the public option panel can fail only through one or more of the following local mechanisms:*

- (1) *the quote vector violates static no-arbitrage or calendar convex-order admissibility;*
- (2) *the public response is not locally injective on the source gauge quotient, equivalently the observability modulus vanishes;*
- (3) *basket or carrier design leaves a nontrivial covariance or synchronization fiber, visible through rank loss in $W_{\mathcal{E}}$ or degeneracy of Γ_{car} ;*
- (4) *the Perron reduced-resolvent norm diverges, in particular through spectral-gap collapse or non-normal conditioning, so the dominant synchronization mode is no longer a smooth coordinate;*
- (5) *the maturity-response dictionary loses rank, so rough or Volterra source coordinates cannot be separated by the admitted maturity panel;*
- (6) *the public payoff module is incomplete for the target claim, so the residual price fiber has nonzero width;*
- (7) *the recovered finite-dimensional state violates the martingale drift restrictions for admitted option coordinates.*

Conversely, away from these failure strata, and under correct source-family specification, the equivalence theorem gives local gauge-class recovery, stability controlled by $\mathbf{m}_t^{-1}$, residual-free pricing for publicly completed claims, and dynamically admissible recovered prices subject to the martingale equations.

Proof. Static and dynamic no-arbitrage are necessary before any risk-neutral interpretation exists. Local source recovery is equivalent to local injectivity of the response on the gauge quotient; in finite dimension, stable recovery requires positive quotient singular value. Basket and carrier rank failures produce explicit kernels in the joint-observation block. Perron reduced-resolvent blow-up invalidates differentiable spectral coordinates; spectral-gap collapse is one mechanism, and non-normal conditioning is another. Maturity-rank collapse creates kernel directions in the Volterra dictionary. Public payoff incompleteness is the existence of a public-annihilating residual direction that moves the target claim. The martingale condition is the dynamic no-arbitrage restriction after recovery. These mechanisms are the local ways in which one clause of Theorem 8.22 can fail. $\square$

Definition 8.26 (Singularly regular observation map). Let $F : U \subset \mathbb{R}^d \rightarrow \mathbb{R}^M$ be an analytic finite-panel response map on a gauge slice. The map is m -regular at u_0 on a closed cone $\mathcal{C} \subset \mathbb{R}^d$ if there are $c > 0$ and $r > 0$ such that, for all $u, v \in u_0 + \mathcal{C}$ with $\|u - u_0\|, \|v - u_0\| < r$,

$$\|F(u) - F(v)\| \geq c \|u - v\|^m.$$

The case $m = 1$ is ordinary Lipschitz recovery. The case $m > 1$ is singular but still identifiable on the cone.

Theorem 8.27 (Hölder recovery across analytic seams). *If F is m -regular at u_0 on the gauge-normalized cone $\mathcal{C}$, then F is locally injective on that cone and its inverse satisfies*

$$\|u - v\| \leq c^{-1/m} \|F(u) - F(v)\|^{1/m}.$$

Consequently, let $y = F(u)$, let $\|\hat{y} - y\| \leq \delta$, and let a reconstruction $\hat{u} \in u_0 + \mathcal{C}$ satisfy

$$\|F(\hat{u}) - \hat{y}\| = o(\delta).$$

Then the local recovery bound

$$\|\hat{u} - u\| \leq c^{-1/m} \delta^{1/m} + o(\delta^{1/m})$$

holds on the same local cone.

Proof. If $F(u) = F(v)$ for two points in the cone, the defining inequality forces $u = v$. Thus F is locally injective on that cone. The same inequality rearranges to the displayed inverse bound. For the noisy estimate,

$$\|F(\hat{u}) - F(u)\| \leq \|F(\hat{u}) - \hat{y}\| + \|\hat{y} - y\| = \delta + o(\delta).$$

Applying the inverse bound gives

$$\|\hat{u} - u\| \leq c^{-1/m} (\delta + o(\delta))^{1/m} = c^{-1/m} \delta^{1/m} + o(\delta^{1/m}).$$

A nonzero projection residual would enter through the first term in the preceding triangle inequality. $\square$

Definition 8.28 (Pseudo-true recovered state). Let $F : U \subset \mathbb{R}^d \rightarrow \mathbb{R}^M$ be a finite-panel response map and let $y \in \mathbb{R}^M$ be an arbitrage-free quote vector, not necessarily in $F(U)$. For a positive definite weight W , a pseudo-true recovered state is

$$u^\dagger \in \arg \min_{u \in U} (y - F(u))^\top W (y - F(u)).$$

The model residual is $r^\dagger := y - F(u^\dagger)$.

Theorem 8.29 (Misspecified projection and normal equations). Assume F is C^2 , $u^\dagger$ is an interior minimizer, and $J^\dagger := DF(u^\dagger)$ has full column rank. Then

$$(J^\dagger)^\top W r^\dagger = 0.$$

If

$$(J^\dagger)^\top W J^\dagger - \sum_{a=1}^M (W r^\dagger)_a D^2 F_a(u^\dagger)$$

is positive definite, then $u^\dagger$ is a strict local projection. Small quote perturbations move $u^\dagger$ differentiably with derivative

$$du^\dagger = \left[(J^\dagger)^\top W J^\dagger - \sum_{a=1}^M (W r^\dagger)_a D^2 F_a(u^\dagger) \right]^{-1} (J^\dagger)^\top W dy.$$

When $r^\dagger = 0$, this reduces to the Gauss–Newton inverse.

Proof. For

$$Q(u; y) = (y - F(u))^\top W (y - F(u)),$$

the first variation is

$$D_u Q(u; y)[h] = -2h^\top DF(u)^\top W (y - F(u)).$$

At an interior minimizer this gives the normal equation $(J^\dagger)^\top W r^\dagger = 0$. The Hessian in the state variable is the displayed positive-definite matrix up to the common factor two, so the second-order condition gives a strict local projection. Applying the implicit-function theorem to the normal equation $DF(u)^\top W (y - F(u)) = 0$ and differentiating with respect to y gives the stated sensitivity formula. $\square$

Definition 8.30 (Weighted finite-panel estimator). Let $F : \mathbb{R}^d \rightarrow \mathbb{R}^M$ be the finite-panel response map on a gauge slice. Observed quotes satisfy

$$Y_n = F(u_0) + n^{-1/2} \varepsilon_n, \quad \varepsilon_n \Rightarrow \mathcal{N}(0, \Omega),$$

where Ω is positive definite on the quote coordinates. For a positive definite weight W , define

$$\hat{u}_n \in \arg \min_{u \in U} (Y_n - F(u))^\top W (Y_n - F(u)).$$

Theorem 8.31 (Asymptotic covariance of recovered states). *Assume F is C^2 , $J := DF(u_0)$ has full column rank, u_0 is an interior point of the gauge slice, and the estimator in Definition 8.30 is locally consistent. Then*

$$\sqrt{n}(\hat{u}_n - u_0) \Rightarrow \mathcal{N}\left(0, (J^\top W J)^{-1} J^\top W \Omega W J (J^\top W J)^{-1}\right).$$

The efficient weight is $W = \Omega^{-1}$, giving covariance

$$\mathcal{I}(u_0)^{-1} := (J^\top \Omega^{-1} J)^{-1}.$$

This is the local Cramér–Rao lower bound for linear functionals in the Gaussian quote experiment generated by the same panel.

Proof. The first-order condition is

$$DF(\hat{u}_n)^\top W(Y_n - F(\hat{u}_n)) = 0.$$

Local consistency and the expansion $Y_n = F(u_0) + n^{-1/2}\varepsilon_n$ give

$$J^\top W \varepsilon_n - J^\top W J \sqrt{n}(\hat{u}_n - u_0) + o_p(1) = 0.$$

Thus

$$\sqrt{n}(\hat{u}_n - u_0) = (J^\top W J)^{-1} J^\top W \varepsilon_n + o_p(1),$$

and the stated covariance follows from $\varepsilon_n \Rightarrow \mathcal{N}(0, \Omega)$. Minimizing this Loewner covariance over positive definite weights gives the generalized least-squares weight $W = \Omega^{-1}$, hence the inverse Fisher information of the local Gaussian shift experiment. $\square$

Corollary 8.32 (Quote-design criterion). *For a candidate finite quote panel with Jacobian J_P and quote-noise covariance Ω_P , the local information matrix is*

$$\mathcal{I}_P := J_P^\top \Omega_P^{-1} J_P.$$

A panel is locally well designed for recovery if $\lambda_{\min}(\mathcal{I}_P)$ is bounded away from zero. For a scalar functional $g(u)$, the local asymptotic variance is

$$\nabla g(u_0)^\top \mathcal{I}_P^{-1} \nabla g(u_0).$$

Thus the theory identifies which strikes, maturities, carriers, and baskets improve recovery of Perron stress, roughness, or directed-route coefficients.

Proof. Theorem 8.31 gives asymptotic covariance $(J_P^\top \Omega_P^{-1} J_P)^{-1}$ for the efficiently weighted local recovery estimator. Hence $\lambda_{\min}(\mathcal{I}_P) > 0$ is the finite-panel information condition, and a lower bound on this eigenvalue prevents loss of quotient directions. The scalar-functional variance follows from the delta method applied to g . $\square$

9. ADAPTED RECOVERY OF THE RISK-NEUTRAL STATE PROCESS

Dynamic observability follows the recovered-state construction of [25]. The preceding theorems are pointwise in t . A dynamic pricing model also requires the recovered state to be adapted through time. Pointwise inverses assemble into an observable $\mathbb{Q}$ -state process only under measurable charts and uniform quotient conditioning.

Definition 9.1 (Local recovery operator). On a recovery neighborhood at date t , define

$$\mathcal{I}_t : W_t \cap F_t(V_t) \rightarrow \bar{\mathcal{X}}_t, \quad \mathcal{I}_t(y) := \bar{F}_t^{-1}(y),$$

where $F_t = \mathcal{R}_t \circ \psi_t$. The recovered canonical state is

$$\hat{\Xi}_t^{\mathbb{Q}} := \mathcal{I}_t(\mathcal{M}_t^{\text{opt}}).$$

Assumption 9.2 (Uniform recovery window). On a time interval $[0, T_0]$, there are $\mathcal{H}_t$ -measurable recovery events Ω_t^{rec} on which the realized option manifold remains in local recovery neighborhoods and the realized response maps satisfy the rank conditions with deterministic lower bounds. On these events the recovery charts can be chosen so that:

- (1) $(t, y) \mapsto \mathcal{I}_t(y)$ is Borel measurable;
- (2) the inverse response norms satisfy $\sup_{t \leq T_0} \|L_t\| < \infty$;

(3) the Perron gaps, Perron reduced-resolvent constants, and carrier Gram lower bounds satisfy

$$\inf_{t \leq T_0} \gamma_P(t) > 0, \quad \sup_{t \leq T_0} \kappa_P(t) < \infty, \quad \inf_{t \leq T_0} \kappa_{\text{car}}(t) > 0;$$

(4) the diagonal maturity, directed-Volterra, and synchronization response blocks have uniformly bounded left inverses; augmented residual-instrument blocks, when used, satisfy the same condition.

Theorem 9.3 (Adapted dynamic recovery). *Suppose Assumption 9.2 holds. Then the recovered canonical state process*

$$\widehat{\Xi}_t^{\mathbb{Q}} = \mathcal{I}_t(\mathcal{M}_t^{\text{opt}})$$

is $\mathcal{H}_t$ -measurable on Ω_t^{rec} for every t . If the public option manifold is $(\mathcal{H}_t)$ -adapted, then $(\widehat{\Xi}_t^{\mathbb{Q}})_{t \leq T_0}$ is an adapted canonical Volterra–Perron state process.

Proof. The public option manifold $\mathcal{M}_t^{\text{opt}}$ is an $\mathcal{H}_t$ -measurable quote object. Assumption 9.2 gives a Borel recovery map on the recovery event at date t . The composition $\mathcal{I}_t(\mathcal{M}_t^{\text{opt}})$ is therefore $\mathcal{H}_t$ -measurable on Ω_t^{rec} . The argument is pointwise in t , and the adaptedness statement follows for the recovered process on the interval. $\square$

Theorem 9.4 (Pathwise stability of recovered states). *Under Assumption 9.2, there is a constant $C_* > 0$ such that for any two public option-manifold paths $M, \widetilde{M}$ remaining in the same recovery neighborhoods,*

$$\sup_{t \leq T_0} \text{dist}_{\widetilde{\mathcal{X}}_t}(\mathcal{I}_t(M_t), \mathcal{I}_t(\widetilde{M}_t)) \leq C_* \sup_{t \leq T_0} \|M_t - \widetilde{M}_t\|_{\mathcal{Y}_t}.$$

If $t \mapsto \mathcal{I}_t$ is locally continuous and $t \mapsto \mathcal{M}_t^{\text{opt}}$ is càdlàg in the quote-space topology, then $t \mapsto \widehat{\Xi}_t^{\mathbb{Q}}$ is càdlàg in the canonical state topology.

Proof. Pointwise stability gives a Lipschitz constant controlled by the inverse response norm, Perron reduced-resolvent bound, carrier Gram lower bound, and diagonal response-block inverse norms. The uniform recovery-window assumption bounds these quantities on $[0, T_0]$; taking the supremum over t gives the displayed estimate. The càdlàg statement follows by composing a càdlàg option-manifold path with locally continuous recovery maps. $\square$

Definition 9.5 (Recovery atlas). A recovery atlas is a family of gauge-normalized charts $(U_\alpha, F_\alpha, \mathcal{I}_\alpha)_{\alpha \in A}$ such that

$$F_\alpha : U_\alpha \rightarrow \mathcal{Y}_t$$

is locally injective on the canonical quotient, $\mathcal{I}_\alpha$ is its inverse on the model image, and on overlaps the inverses agree up to the prescribed gauge transition:

$$\mathcal{I}_\beta(y) = g_{\alpha\beta}(\mathcal{I}_\alpha(y)), \quad y \in F_\alpha(U_\alpha) \cap F_\beta(U_\beta).$$

The atlas is uniformly observable on a region $\mathcal{O}$ if all local inverse constants are bounded there by a common constant.

Theorem 9.6 (Patching local recovery). *Let $\mathcal{M}_t^{\text{opt}}$ be an adapted option-manifold path taking values in the image of a uniformly observable recovery atlas on a region $\mathcal{O}$. Suppose the chart transitions are Borel. Then the local recovered states patch into a well-defined adapted canonical state process on the stopping interval $[0, \tau_{\text{loss}})$, where*

$$\tau_{\text{loss}} := \inf\{t : \mathcal{M}_t^{\text{opt}} \notin \mathcal{O} \text{ or an observability boundary is hit}\}.$$

The observability boundary may include

$$\gamma_P(t) = 0, \quad \lambda_{\min}(\Gamma_{\text{car},t}) = 0, \quad \mathfrak{D}_t(\tau) = 0,$$

as well as active-edge transitions $V_{ij,t} = 0$ for screened routes whose stratum is being held fixed.

Proof. On each chart domain, adaptedness follows from the local inverse theorem. On overlaps, gauge-compatible transition maps identify the same canonical state. The first exit time from the observable region is a stopping time when the boundary functions are adapted and the recovery region is measurable. Uniform inverse bounds give pathwise stability before exit. $\square$

Remark 9.7. The dynamic recovery theorem is an adapted inverse-map result, with no observation-noise filtering model. Statistical filtering would require an additional observation-error model for finite quotes, bid-ask frictions, and microstructure smoothing. The map constructed here is the quotient inverse that such a filter would approximate.

Definition 9.8 (Recovered pricing semigroup). Let U_t denote a finite-dimensional recovered quotient state on a recovery window. For a payoff class $\mathcal{D}$ and maturity increment $\tau \geq 0$, define

$$P_{t,t+\tau}f(U_t) := \mathbb{E}^{\mathbb{Q}} \left[\frac{B_t}{B_{t+\tau}} f(U_{t+\tau}) \middle| \mathcal{H}_t \right].$$

The option manifold prices a core $\mathcal{D}$ if these values are observable or statically spanned for every $f \in \mathcal{D}$ and all sufficiently small τ .

Theorem 9.9 (Generator identification from short maturities). *Assume U is Markov under $\mathbb{Q}$ on the recovery window, and let $\mathcal{A}_t$ be its risk-neutral generator on a core $\mathcal{D}$. If the recovered option manifold identifies $P_{t,t+\tau}f(U_t)$ for every $f \in \mathcal{D}$ and all sufficiently small $\tau > 0$, then*

$$\mathcal{A}_t f(U_t) - r_t f(U_t) = \lim_{\tau \downarrow 0} \frac{P_{t,t+\tau}f(U_t) - f(U_t)}{\tau}$$

for every $f \in \mathcal{D}$ for which the limit exists. If

$$dU_t = b(t, U_t) dt + \Sigma(t, U_t) dW_t^{\mathbb{Q}}$$

is an Itô diffusion, then prices of coordinate functions and their quadratic products identify first the discounted generator

$$L_t f(U_t) := \lim_{\tau \downarrow 0} \frac{P_{t,t+\tau}f(U_t) - f(U_t)}{\tau} = \mathcal{A}_t f(U_t) - r_t f(U_t),$$

and hence the drift and covariance coefficients by

$$b_i(t, U_t) = L_t u_i(U_t) + r_t u_i(U_t),$$

and

$$a_{ij}(t, U_t) = L_t(u_i u_j)(U_t) + r_t u_i(U_t) u_j(U_t) - u_i(U_t) b_j(t, U_t) - u_j(U_t) b_i(t, U_t), \quad a = \Sigma \Sigma^\top.$$

Proof. On the core $\mathcal{D}$, the recovered short-maturity prices identify the discounted pricing semigroup to first order:

$$P_{t,t+\tau}f(U_t) = f(U_t) + \tau(\mathcal{A}_t f(U_t) - r_t f(U_t)) + o(\tau).$$

The displayed limit is therefore the generator of the discounted semigroup, so the option manifold identifies $L_t = \mathcal{A}_t - r_t$ on the core, and adding back $r_t f$ recovers $\mathcal{A}_t f$. For the diffusion statement, apply

$$\mathcal{A}_t f = b^\top \nabla f + \frac{1}{2} \text{Tr}(a \nabla^2 f)$$

to $f(u) = u_i$ and $f(u) = u_i u_j$, then replace $\mathcal{A}_t$ by $L_t + r_t$ in the identified quantities. $\square$

Remark 9.10 (Ideal generator recovery). Theorem 9.9 is an ideal short-maturity result. It requires the public manifold, or an admitted static span of it, to price the coordinate functions and quadratic products in the chosen core. Ordinary listed vanillas do not automatically contain these state claims. When the core is only approximated by a finite quote panel, the finite state-law experiment below should be read as an adapted propagation and innovation diagnostic, not as exact identification of the infinitesimal generator.

Definition 9.11 (Finite adapted state-law experiment). Let $0 = t_0 < t_1 < \dots < t_n \leq T_0$ be observation dates inside a recovery window, and let

$$U_k := \widehat{\Xi}_{t_k}^{\mathbb{Q}} \in \mathbb{R}^d$$

be a gauge-normalized recovered coordinate vector in a fixed chart. Write $\Delta_k = t_{k+1} - t_k$, $\Delta U_k = U_{k+1} - U_k$, and $\mathcal{G}_k = \mathcal{H}_{t_k} \vee \sigma(U_0, \dots, U_k)$. A finite adapted state-law experiment specifies, for each k , a coefficient estimate

$$\widehat{b}_k^-(u) = \widehat{a}_k + \widehat{B}_k u$$

and positive semidefinite innovation matrix $\hat{A}_k^-$, both $\mathcal{G}_k$ -measurable and estimated only from transitions (U_j, U_{j+1}) with $j < k$. The one-step state-law innovation is

$$\varepsilon_{k+1} := U_{k+1} - U_k - \Delta_k \hat{b}_k^-(U_k),$$

and the persistence innovation is $\varepsilon_{k+1}^0 := \Delta U_k$.

Definition 9.12 (Phase-local diffusion and Gaussian seam energy). Let $\phi_k = \phi(U_k)$ be any hard or smooth phase label generated by the quotient phase coordinates of Section 5. Let $\hat{A}_k^{\text{glob}}$ be a shrunk covariance estimate of the scaled past innovations

$$\frac{\varepsilon_{j+1}}{\sqrt{\Delta_j}}, \quad j < k,$$

and let $\hat{A}_k^\phi$ be the analogous estimate restricted to past indices $j < k$ with $\phi_j = \phi_k$, whenever enough such observations exist. Otherwise set $\hat{A}_k^\phi = \hat{A}_k^{\text{glob}}$ and record that the phase-local covariance was not separately identified. With ridge parameter $\eta_k \geq 0$, define the global and phase-local Gaussian seam energies

$$\mathfrak{E}_{k+1}^{\text{glob}} = \frac{1}{d} \left(\frac{\varepsilon_{k+1}}{\sqrt{\Delta_k}} \right)^\top \left(\hat{A}_k^{\text{glob}} + \eta_k I \right)^\dagger \left(\frac{\varepsilon_{k+1}}{\sqrt{\Delta_k}} \right),$$

and

$$\mathfrak{E}_{k+1}^\phi = \frac{1}{d} \left(\frac{\varepsilon_{k+1}}{\sqrt{\Delta_k}} \right)^\top \left(\hat{A}_k^\phi + \eta_k I \right)^\dagger \left(\frac{\varepsilon_{k+1}}{\sqrt{\Delta_k}} \right).$$

The same definitions with ε_{k+1}^0 give the corresponding persistence seam energies.

Proposition 9.13 (Adaptedness of finite state-law diagnostics). *In the finite adapted state-law experiment, the fitted drift $\hat{b}_k^-$, global covariance $\hat{A}_k^{\text{glob}}$, phase-local covariance $\hat{A}_k^\phi$, ridge η_k , and forecast*

$$U_k + \Delta_k \hat{b}_k^-(U_k)$$

are $\mathcal{G}_k$ -measurable. Therefore the one-step propagation error, innovation scale, and seam energies are prequential diagnostics: they evaluate the realized next recovered state against an object formed before U_{k+1} is observed.

Proof. All coefficient estimates and covariance estimates are functions only of $\mathcal{G}_k$ -measurable past transitions and of the current recovered state U_k . The phase label ϕ_k is a measurable function of U_k . Hence the fitted drift, covariance matrices, ridge, and forecast are $\mathcal{G}_k$ -measurable. The diagnostics are then obtained by comparing these known-at- t_k objects with U_{k+1} . $\square$

Theorem 9.14 (Finite state-law consistency). *Suppose on a recovery window the recovered state satisfies the Itô diffusion*

$$dU_t = b(t, U_t) dt + \Sigma(t, U_t) dW_t^\mathbb{Q}, \quad A(t, U_t) = \Sigma(t, U_t) \Sigma(t, U_t)^\top,$$

with locally Lipschitz coefficients. Let $h_k = \max_{j < k} \Delta_j \rightarrow 0$, and let the regression and covariance estimators in Definition 9.11 be consistent local estimators of $b(t_k, U_{t_k})$ and $A(t_k, U_{t_k})$ on the chosen chart or phase stratum. Then

$$\frac{\varepsilon_{k+1}}{\sqrt{\Delta_k}} = \Sigma(t_k, U_{t_k}) Z_{k+1} + o_{\mathbb{Q}}(1), \quad Z_{k+1} \sim N(0, I)$$

conditionally to first order, and the seam energy in Definition 9.12 is a finite-information innovation statistic for the recovered $\mathbb{Q}$ -state. If the local affine drift is misspecified, the same statistic remains an adapted propagation-reserve diagnostic, not a source-transition estimator for the true law.

Proof. The Euler expansion of the Itô diffusion gives

$$U_{t_{k+1}} - U_{t_k} = b(t_k, U_{t_k}) \Delta_k + \Sigma(t_k, U_{t_k}) (W_{t_{k+1}} - W_{t_k}) + o_{\mathbb{Q}}(\sqrt{\Delta_k})$$

under the stated regularity. Consistency of the drift estimator makes the fitted drift error lower order after multiplication by Δ_k . Consistency of the covariance estimator gives the

displayed covariance-normalized innovation statistic. Without correct local specification, the same construction is still prequential by Proposition 9.13; it is then a reserve or fragility measure rather than a recovered transition kernel. $\square$

Remark 9.15 (Role of the Gaussian seam metric). The seam energy is the dynamic analogue of the finite Gaussian source-family tests used in the recovery layer. Raw Euclidean propagation errors depend on the chosen coordinate scale. The covariance-normalized innovation measures the error in the finite information metric of the recovered chart. Phase-local versions replace a pooled covariance by the covariance of comparable past quotient states; they do not create a new phase by thresholding future errors.

Remark 9.16 (Pricing dynamics versus observed state dynamics). The consistency statement above is a $\mathbb{Q}$ -dynamics statement. A calendar time series of option-implied recovered states is usually observed under the statistical data-generating measure, even though each state is a risk-neutral object recovered from option prices. Such a time series therefore estimates observation-measure dynamics of $\mathbb{Q}$ -implied coordinates unless an additional change-of-measure or state-space transfer map is imposed. The finite prequential diagnostics remain valid as adapted stability, propagation, and fragility diagnostics for the observed option manifold; they are not direct estimates of the risk-neutral pricing generator without that extra transfer map.

10. MARTINGALE CONSISTENCY OF RECOVERED RISK-NEUTRAL DYNAMICS

Recovery turns the option manifold into a canonical state process. Risk-neutral pricing imposes a separate condition. Once public prices are written as functionals of the recovered state, their discounted dynamics must be local martingales. The resulting drift equations are the dynamic no-arbitrage restrictions on the Volterra–Perron coordinates.

Definition 10.1 (Recovered price functional). In this section a denotes a fixed calendar contract, such as (u, K, T, κ) , rather than a rolling fixed- τ surface point. Let $a \in \mathcal{I}_t$ be an admitted option coordinate. On a recovery neighborhood, define

$$\Phi_a(t, \bar{\Theta}) := C_t^a \quad \text{whenever} \quad \bar{\Theta} = \mathcal{I}_t(\mathcal{M}_t^{\text{opt}}).$$

Equivalently, Φ_a is the a -coordinate of the observable response map expressed in canonical recovered-state coordinates.

Remark 10.2. If the same restriction is written in rolling surface coordinates $(u, \tau, m, \kappa, \chi)$, the generator must include the roll of $\tau = T - t$ and, when m is defined relative to a stochastic forward, the Itô generator of m . In the simplest fixed strike parametrization this adds the familiar convective term $-\partial_\tau$ to the surface PDE.

Assumption 10.3 (Finite-dimensional semimartingale chart). On a recovery window, the canonical quotient state admits local coordinates $U_t \in \mathbb{R}^d$ and has risk-neutral semimartingale dynamics

$$dU_t = b_t dt + \Sigma_t dW_t^{\mathbb{Q}},$$

where $b_t \in \mathbb{R}^d$, $\Sigma_t \in \mathbb{R}^{d \times q}$, and the recovered price functionals $\Phi_a(t, u)$ are $C^{1,2}$ in (t, u) on the coordinate neighborhood.

Theorem 10.4 (Recovered-state martingale restriction). *Suppose Assumption 10.3 holds and, for every admitted public option coordinate a , the discounted price process*

$$e^{-\int_0^t r_s ds} \Phi_a(t, U_t)$$

is a local $\mathbb{Q}$ -martingale. Then, along the recovered state path,

$$\partial_t \Phi_a(t, U_t) + D\Phi_a(t, U_t)[b_t] + \frac{1}{2} \text{Tr} \left(\Sigma_t^* D^2 \Phi_a(t, U_t) \Sigma_t \right) - r_t \Phi_a(t, U_t) = 0$$

for every admitted option coordinate a , up to indistinguishability.

Proof. Apply Itô's formula to $e^{-\int_0^t r_s ds} \Phi_a(t, U_t)$. The finite-variation drift is the discount factor times

$$\partial_t \Phi_a + D\Phi_a[b_t] + \frac{1}{2} \text{Tr}(\Sigma_t^* D^2 \Phi_a \Sigma_t) - r_t \Phi_a.$$

A local martingale has zero finite-variation drift. This gives the displayed restriction. $\square$

Corollary 10.5 (Dynamic no-arbitrage closure). *Conversely, suppose the drift restriction in Theorem 10.4 holds for all admitted coordinates and the stochastic integrals are well defined. Then the generated discounted public option price coordinates are local $\mathbb{Q}$ -martingales. Hence the recovered state dynamics define a dynamically arbitrage-free public option manifold on the admitted payoff span, subject to the usual true-martingale integrability conditions.*

Proof. The same Itô calculation shows that the finite-variation part of each discounted coordinate vanishes. The remaining term is a stochastic integral and hence a local martingale. A positive linear conditional-expectation representation on the admitted span gives local no-arbitrage; true martingale integrability upgrades local prices to genuine conditional-expectation prices. $\square$

Definition 10.6 (Recovered Musiela option surface). Let U_t be a finite-dimensional recovered state chart. For underlier u , define the discounted rolling surface coordinate in pricing chart $\chi \in \{\mathbb{B}, \mathbb{N}\}$ by

$$\mathcal{C}^{u,\chi}(t, U_t, \tau, m^\chi) := B_t^{-1} C_t^u(T = t + \tau, K = K^\chi(F_t^u, m^\chi)),$$

where $\tau > 0$ is time to maturity,

$$K^{\mathbb{B}}(F, m) = F e^m, \quad K^{\mathbb{N}}(F, m) = F + m.$$

For a fixed calendar contract (u, K, T) , set

$$\tau_t = T - t, \quad m_t^{\mathbb{B}} = \log(K/F_t^u), \quad m_t^{\mathbb{N}} = K - F_t^u.$$

Thus the Black roll is written in log-moneyness, whereas the normal-forward roll is written in cash-forward moneyness. A standardized normal coordinate $(K - F_t^u)/\sqrt{v_{u,t}^{\mathbb{N}}(\tau)}$ may be used instead, but then the generator also contains the roll of the normal variance scale. Let $\mathcal{L}_{U,m^\chi}$ denote the generator of the joint state (U_t, m_t^χ) .

Theorem 10.7 (Surface roll restriction). *Suppose the discounted calendar-contract price*

$$B_t^{-1} C_t^u(K, T) = \mathcal{C}^{u,\chi}(t, U_t, \tau_t, m_t^\chi)$$

is a local $\mathbb{Q}$ -martingale and the joint state (U_t, m_t^χ) has generator $\mathcal{L}_{U,m^\chi}$. Then, along the recovered state path,

$$\partial_t \mathcal{C}^{u,\chi} - \partial_\tau \mathcal{C}^{u,\chi} + \mathcal{L}_{U,m^\chi} \mathcal{C}^{u,\chi} = 0.$$

In undiscounted coordinates the right-hand side is $r_t \mathcal{C}^{u,\chi}$. If m^χ is held fixed as a surface coordinate rather than induced by a fixed calendar strike, the generator must be changed accordingly; an additional chart-specific moneyness-roll term enters.

Proof. Apply Itô's formula to $\mathcal{C}^{u,\chi}(t, U_t, T - t, m_t^\chi)$. The deterministic roll of τ_t contributes $-\partial_\tau \mathcal{C}^{u,\chi}$, while the joint diffusion and drift of (U_t, m_t^χ) contribute $\mathcal{L}_{U,m^\chi} \mathcal{C}^{u,\chi}$. The finite-variation part of a discounted local martingale must vanish. $\square$

Definition 10.8 (Dynamically admissible residual selector). A residual selector θ is dynamically admissible on a recovery window if the selected price functionals

$$\Phi_X^\theta(t, U_t) := \Pi_t^\theta(X)$$

for the selected latent-stress claims satisfy the same discounted local-martingale drift restriction as public option coordinates, whenever they are included in the priced claim family.

Proposition 10.9 (Residual selectors are constrained but not eliminated). *Dynamic martingale consistency constrains the drift of selected residual price functionals, while a separate completion condition is needed to collapse the residual interval to a single selector. If two positive public-equivalent selector density families generate selected price functionals satisfying the drift*

restriction and agree on public option coordinates, they remain observationally equivalent on the public manifold while assigning different prices to residual latent-stress claims.

Proof. Agreement on public option coordinates follows from public equivalence. The martingale restriction says that each selected discounted price process has zero drift. This gives dynamic consistency for an already selected extension. Completeness of the tradeable span requires a separate condition. If the two selector density families differ on a residual direction correlated with the latent claim, the public non-identification calculation gives different selected prices for that claim while preserving all public prices. $\square$

Definition 10.10 (Process-consistent public-equivalent kernel). A horizon-wise public-equivalent family $Z_{t,T}$ is process-consistent if there is a positive $\mathbb{Q}$ -martingale density process $D^\theta = (D_s^\theta)_{s \geq 0}$, $D_0^\theta = 1$, such that

$$Z_{t,T}^\theta = \frac{D_T^\theta}{D_t^\theta}.$$

It is dynamically public-equivalent on the admitted manifold if, for every date t and every admitted public payoff X_a with maturity $T(a) > t$,

$$\mathbb{E}^\mathbb{Q} \left[\frac{D_{T(a)}^\theta}{D_t^\theta} X_a \middle| \mathcal{H}_t \right] = \mathbb{E}^\mathbb{Q} [X_a \mid \mathcal{H}_t].$$

Theorem 10.11 (Martingale residual densities preserve public prices). *Let D^θ be a positive martingale density process satisfying the dynamic public-equivalence identities in Definition 10.10. Then D^θ defines a dynamically public-equivalent extension of the public option manifold. For any latent claim X whose discounted payoff has nonzero conditional covariance with D_T^θ/D_t^θ ,*

$$\mathbb{E}^\mathbb{Q} \left[\frac{D_T^\theta}{D_t^\theta} X \middle| \mathcal{H}_t \right] \neq \mathbb{E}^\mathbb{Q} [X \mid \mathcal{H}_t],$$

the density process changes the selected residual price without changing admitted public option prices.

Proof. The first statement is the defining preservation identity applied at every date and maturity. For a latent claim, the price change is

$$\mathbb{E}^\mathbb{Q} \left[\left(\frac{D_T^\theta}{D_t^\theta} - 1 \right) X \middle| \mathcal{H}_t \right].$$

Because the density ratio has conditional mean one, this term is the conditional covariance of the payoff with D_T^θ/D_t^θ . Nonzero conditional covariance is therefore exactly the condition that the selected residual price differs from the baseline conditional price on a set of positive probability. $\square$

Remark 10.12. The drift restriction is the option-manifold analogue of an HJM restriction. The state may be parametrized by Volterra kernels, Perron spectral coordinates, rough maturity coordinates, and residual selectors, but once these coordinates generate public option prices under $\mathbb{Q}$, their joint drift and diffusion cannot be chosen independently.

11. CANONICAL FINITE BENCHMARK MODEL

The preceding sections establish when a risk-neutral Volterra–Perron source state is recoverable and which residual prices remain undetermined. The finite benchmark fixes a closed chart: the recovered state is a finite vector, the option formula is explicit, the no-arbitrage equation is a state PDE, native Greeks are derivatives of the same state, and hedge incompleteness has a computable residual.

The benchmark uses Bachelier coordinates. The Bachelier choice is a normal-forward normalization rather than an economic restriction on the theory. It avoids basket-lognormal approximations and makes the synchronization covariance object exact. Lognormal, displaced, or stochastic-volatility variants can be built on the same recovered state, but the finite benchmark is the smallest closed model in which the Volterra–Perron quotient becomes a pricing equation.

Definition 11.1 (Finite Bachelier–Volterra–Perron state). Fix N underliers, a maturity horizon $\bar{\tau}$, and an active set of carrier baskets $\mathcal{B} = \{w^1, \dots, w^M\} \subset \mathbb{R}^N$, including single names, indices, and any admitted carrier portfolios. A finite Bachelier–Volterra–Perron state at date t consists of

$$U_t = (F_t, \nu_t, \mathfrak{H}_t, G_t, \rho_t, e_t, \ell_t, \varpi_t, \theta_t),$$

where $F_t \in \mathbb{R}^N$ is the vector of forward levels, $\nu_t \in (0, \infty)^N$ is a diagonal idiosyncratic variance scale, $\mathfrak{H}_t \in (-1/2, \infty)^Q$ is the finite rough maturity-scaling vector, G_t is the finite directed route matrix, (ρ_t, e_t, ℓ_t) is the Perron spectral triple with $\ell_t(e_t) = 1$, ϖ_t is the synchronized stress coordinate, and θ_t is an admissible residual selector. The usual rough-volatility benchmark takes $\mathfrak{H}_{q,t} \in (0, 1/2]$, but the finite normal-forward covariance formula only requires $2\mathfrak{H}_{q,t} + 1 > 0$ for variance growth at the origin.

For $0 < \tau \leq \bar{\tau}$, the conditional covariance matrix of forward changes is

$$\Gamma_t(\tau; U_t) = D_t(\tau) + \sum_{q=1}^Q \tau^{2\mathfrak{H}_{q,t}+1} B_{q,t} B_{q,t}^\top + a_\varpi(\tau; U_t) c_t c_t^\top + \Gamma_t^{\text{res}}(\tau; \theta_t),$$

where

$$\begin{aligned} D_t(\tau) &= \tau \operatorname{diag}(\nu_t), & c_t &= \operatorname{diag}(\sqrt{\nu_t}) e_t, \\ a_\varpi(\tau; U_t) &= \rho_t \exp(\varpi_t) \tau^{2\bar{\mathfrak{H}}_t+1}, & \bar{\mathfrak{H}}_t &= \frac{\sum_q \omega_q \mathfrak{H}_{q,t}}{\sum_q \omega_q}, \quad \omega_q > 0, \end{aligned}$$

and the matrices $B_{q,t} = B_q(G_t)$ are fixed differentiable functions of the directed route coordinates. The residual block is required to be symmetric and small enough that

$$\Gamma_t(\tau; U_t) \succeq 0, \quad \Gamma_t(\tau_2; U_t) - \Gamma_t(\tau_1; U_t) \succeq 0 \quad (0 < \tau_1 < \tau_2 \leq \bar{\tau})$$

on the admitted maturity interval. Equivalently, in a smooth maturity chart $\partial_\tau \Gamma_t(\tau; U_t) \succeq 0$. The first condition gives slice-wise covariance positivity; the second is the normal-forward calendar condition induced by convex order. The residual-free benchmark is obtained by setting $\Gamma_t^{\text{res}} \equiv 0$.

Remark 11.2 (Gauge and positivity). The gauge is fixed by e_t in the positive cone and $\ell_t(e_t) = 1$; the scalar size of the Perron mode is carried by $\exp(\varpi_t) \rho_t$. This keeps the covariance contribution positive while preserving a centered stress coordinate ϖ_t . Residual selectors do not have to be positive semidefinite term by term; they must keep the total covariance matrix positive semidefinite on the admitted panel. This is the finite counterpart of the residual-envelope positivity condition.

Remark 11.3 (Covariance-projected residual selectors). The residual term $\Gamma_t^{\text{res}}(\tau; \theta_t)$ is a finite normal-forward specialization of the general residual-kernel selector. It represents only the part of the selected residual pricing rule that can be expressed as an admissible covariance correction in a Bachelier chart. General public-equivalent density selectors may also change higher moments, tail charges, or non-Gaussian residual prices without admitting a covariance-only representation. Thus the finite benchmark is a closed pricing model for the covariance-projected branch of the theory, not a replacement for the full residual-density envelope.

Definition 11.4 (Benchmark option coordinates). For a basket $w \in \mathcal{B}$, define

$$F_t^w = w^\top F_t, \quad V_t^w(\tau; U_t) = w^\top \Gamma_t(\tau; U_t) w.$$

Let $P(t, T) > 0$ be the public zero-coupon price and let $\mathbb{Q}^T$ denote the T -forward measure on the admitted payoff span. The vector F_t is the date- t vector of T -forward levels for the contract maturity under consideration; the maturity argument is suppressed only to lighten notation. Let Φ and φ denote the standard normal distribution and density, and put

$$d_t^w(K, T) = \frac{F_t^w - K}{\sqrt{V_t^w(T - t; U_t)}}.$$

For $\kappa = 1$ for a call and $\kappa = -1$ for a put, the finite benchmark price is

$$C_t^\kappa(w, K, T; U_t) = P(t, T) \left[\kappa (F_t^w - K) \Phi(\kappa d_t^w) + \sqrt{V_t^w(T - t; U_t)} \varphi(d_t^w) \right],$$

with the intrinsic-value convention when $V_t^w = 0$. If discounting is deterministic or $\mathcal{H}_t$ -measurable, this is equivalently the money-market formula with $P(t, T) = \mathbb{E}^\mathbb{Q}[B_t/B_T \mid \mathcal{H}_t]$.

Theorem 11.5 (Canonical finite pricing formula). *Suppose that, conditional on $\mathcal{H}_t$ under the T -forward measure $\mathbb{Q}^T$, every admitted basket forward satisfies*

$$F_T^w - F_t^w \sim N(0, V_t^w(T - t; U_t))$$

jointly across the basket family, with covariance induced by $\Gamma_t(T - t; U_t)$. Then the price in Definition 11.4 is the arbitrage-free date- t price of the Bachelier call or put on $w^\top F_T$. The full finite option panel is the map

$$\mathfrak{P}_t^{\text{fin}} : U_t \mapsto (C_t^\kappa(w, K, T; U_t))_{(w, K, T, \kappa) \in \mathcal{I}_t}.$$

If the Jacobian of this map on the gauge-normalized source quotient has full column rank, the finite panel locally recovers the benchmark state modulo the fixed gauge.

Proof. The conditional Bachelier formula is the expectation of $(\kappa(w^\top F_T - K))_+$ under a one-dimensional normal law with mean F_t^w and variance V_t^w under $\mathbb{Q}^T$, multiplied by $P(t, T)$. Joint Gaussianity ensures that all basket prices come from one positive T -forward conditional law for the admitted forward vector, so the finite panel is arbitrage-free on the basket payoff span. The recovery statement is the finite-dimensional inverse-function theorem applied to the gauge-normalized response map. $\square$

Theorem 11.6 (Recovered-state pricing equation in the finite benchmark). *Assume the benchmark state $U_t \in \mathbb{R}^d$ is Markov under $\mathbb{Q}$ on a recovery window,*

$$dU_t = b^\mathbb{Q}(t, U_t) dt + \Sigma^\mathbb{Q}(t, U_t) dW_t^\mathbb{Q}, \quad a^\mathbb{Q} = \Sigma^\mathbb{Q}(\Sigma^\mathbb{Q})^\top,$$

and let $V(t, u)$ be the selected price of a calendar claim written in benchmark coordinates. Then dynamic no-arbitrage requires

$$\partial_t V(t, u) + b^\mathbb{Q}(t, u)^\top \nabla_u V(t, u) + \frac{1}{2} \text{Tr} \left(a^\mathbb{Q}(t, u) D_u^2 V(t, u) \right) - r_t V(t, u) = 0$$

on the recovery window, with terminal condition given by the claim payoff. The second-order term is the native state-volatility charge; it is the finite benchmark form of quotient gamma pricing.

Proof. This is Theorem 10.4 in the finite Markov chart. Applying Itô's formula to the discounted selected price $e^{-\int_0^t r_s ds} V(t, U_t)$, and setting its finite-variation drift equal to zero, gives the displayed equation. $\square$

Proposition 11.7 (Closed native Greeks in the Bachelier benchmark). *Let z be any differentiable coordinate of the finite state U_t . For an admitted basket option with $V = V_t^w(\tau; U_t) > 0$, $F = F_t^w$, and $d = (F - K)/\sqrt{V}$,*

$$\partial_z C_t^\kappa = P(t, T) \left[\kappa \Phi(\kappa d) \partial_z F + \frac{\varphi(d)}{2\sqrt{V}} \partial_z V \right].$$

In particular, for the Perron stress coordinate,

$$\partial_\omega V_t^w(\tau) = a_\omega(\tau; U_t)(w^\top c_t)^2 + w^\top \partial_\omega \Gamma_t^{\text{res}}(\tau; \theta_t) w,$$

for the Perron spectral radius coordinate,

$$\partial_\rho V_t^w(\tau) = \exp(\varpi_t) \tau^{2\mathfrak{H}_t+1} (w^\top c_t)^2 + w^\top \partial_\rho \Gamma_t^{\text{res}}(\tau; \theta_t) w,$$

and, for a roughness coordinate $\mathfrak{H}_q$ appearing in a route term $\tau^{2\mathfrak{H}_q+1} B_q B_q^\top$,

$$\partial_{\mathfrak{H}_q} V_t^w(\tau) = 2 \log(\tau) \tau^{2\mathfrak{H}_q+1} (w^\top B_q)(B_q^\top w) + \partial_{\mathfrak{H}_q} a_\omega(\tau; U_t)(w^\top c_t)^2 + w^\top \partial_{\mathfrak{H}_q} \Gamma_t^{\text{res}}(\tau; \theta_t) w.$$

Directed-route, carrier, eigenvector, and residual-selector Greeks are obtained by the same formula with $\partial_z V = w^\top (\partial_z \Gamma_t) w$.

Proof. Differentiate the Bachelier formula. The terms containing $\partial_z d$ cancel, leaving $\partial_F C = P(t, T)\kappa\Phi(\kappa d)$ and $\partial_V C = P(t, T)\varphi(d)/(2\sqrt{V})$. The displayed native Greek formulas are the corresponding derivatives of the finite covariance matrix in Definition 11.1. The zero-coupon factor is treated as public and not differentiated in this native state chart; if a rate factor is included as a state coordinate, its bond-price derivative adds a separate term. $\square$

Definition 11.8 (Finite public hedge residual). Let $y = \mathfrak{P}_t^{\text{fin}}(U_t) \in \mathbb{R}^M$ be the finite quote vector and $J = D_U \mathfrak{P}_t^{\text{fin}}(U_t)$ its quotient Jacobian. Let $g_X \in (\mathbb{R}^d)^*$ be the native Greek row of a claim X . A finite public hedge for X is a vector $h \in \mathbb{R}^M$ of quote-coordinate holdings. Its first-order native exposure is $h^\top J$. The minimum weighted first-order hedge residual is

$$\mathcal{R}_t^{\text{hedge}}(X) = \inf_{h \in \mathbb{R}^M} \|g_X - h^\top J\|_{\mathcal{W}_t^{-1}},$$

where $\mathcal{W}_t$ is a positive definite state-coordinate risk weight.

Theorem 11.9 (Finite hedge formula and incompleteness certificate). *Let J and g_X be as in Definition 11.8. If J has full column rank, then $\mathcal{R}_t^{\text{hedge}}(X) = 0$, and the minimum quote-norm hedge solving $J^\top h = g_X^\top$ is*

$$h^* = J(J^\top J)^{-1} g_X^\top.$$

If the admitted quote span is restricted to a submatrix J_{tr} , then

$$\mathcal{R}_{t,\text{tr}}^{\text{hedge}}(X) = \left\| (I - \Pi_{\text{Ran}(\mathcal{W}_t^{-1/2} J_{\text{tr}}^\top)}) \mathcal{W}_t^{-1/2} g_X^\top \right\|_2.$$

Thus the finite benchmark separates recovered hedgeable claims ($\mathcal{R} = 0$), recovered unhedgeable claims ($\mathcal{R} > 0$), and unrecovered claims when J is rank deficient on the quotient.

Proof. The full-rank case is the adjoint Greek equation from Definition 14.7 with $W = I$. The minimum Euclidean-norm solution is the displayed Moore–Penrose solution. With a restricted tradeable quote span, the best hedge is the orthogonal projection of the weighted native Greek onto the weighted row space generated by the tradeable Jacobian. The residual norm is exactly the distance to that row space. $\square$

Remark 11.10 (What the finite benchmark does and does not solve). The benchmark is the finite analogue of a canonical pricing equation. It makes no claim that the full listed-option manifold is Gaussian or Bachelier. It gives a closed reference model in which the Volterra–Perron state, option formula, state PDE, native Greeks, and public hedge residual are all explicit. Smile completion, jump tails, and liquidity effects are outside this finite reference model. The mathematical role of the benchmark is to show that the quotient geometry can be compressed into a tractable no-arbitrage model without losing the residual and carrier distinctions that drive the general theory.

12. GAUSSIAN–VOLTERRA SOURCE-FAMILY RECOVERY

The Gaussian–Volterra branch is a source-family verification of the quotient recovery theorem. The reduced form keeps the rank mechanism visible: temporal singular profiles, scalar seams, and a common Perron-aligned option channel supply the response blocks. Companion A supplies the primitive singular-value and scalar-seam estimates [26].

Assumption 12.1 (Gaussian–Volterra recovery family). The risk-neutral latent volatility system is Gaussian conditional on $\mathcal{H}_t$. Its directed Volterra kernels have nonvanishing diagonal amplitudes and belong to a finite source chart whose coordinates determine:

- (1) the leading singular-value profile of each temporal Volterra route;
- (2) the centered quadratic seam coefficient;
- (3) the common Perron-aligned option channel;
- (4) the index/single-name synchronization response matrix over an open maturity interval.

The chart is gauge-normalized by fixing driver basis, Perron normalization, and common-channel sign as in Definition 3.14.

Assumption 12.2 (Gaussian response rank). The Gaussian response blocks listed below have full rank after quotienting by gauge:

- (1) the temporal singular-value profile block that identifies rough maturity scaling;
- (2) the Volterra route block generated by cross-kernel singular coefficients;
- (3) the common-channel index/carrier response matrix;
- (4) in an augmented residual-instrument manifold, the residual-selector response on the admitted residual-direction chart.

The leading singular-value coefficients in the temporal routes are nonzero, and the seam defect is lower order uniformly on the maturity interval.

Theorem 12.3 (Reduced Gaussian–Volterra recovery). *Under Assumptions 12.1 and 12.2, and assuming a simple well-conditioned Perron eigenvalue and nondegenerate carrier Gram matrix, the Gaussian–Volterra option manifold locally recovers the canonical risk-neutral Volterra–Perron state. If the residual-selector rank block is supplied only by augmented residual instruments, the selector is recovered only in that augmented manifold; otherwise it remains a selected pricing extension.*

Proof. The Gaussian source-family assumptions provide the chart U of Assumption 8.1. The leading singular-value profile and scalar seam identify the temporal and common-channel coordinates in the chart up to the gauge conventions. The Volterra route block identifies the directed cross-kernel coordinates, while the common-channel matrix identifies the Perron synchronization coordinates. The response-rank assumption gives bounded left inverses for the diagonal blocks in Definition 8.12. Finite dimensionality of the source chart turns injectivity of these blocks into bounded inversion on their ranges. A simple Perron eigenvalue with bounded reduced resolvent makes the Perron eigenvalue and eigenvectors differentiable functions of the screened operator through the simple-isolated-eigenvalue perturbation formulas. The carrier Gram condition prevents collapse of the carrier coordinates. Hence the block-rank observability criterion applies, and Theorem 8.4 gives canonical recovery. $\square$

Remark 12.4. The Gaussian companion analysis is used as a verification layer, not as a separate economic claim. Its role is to justify that primitive Gaussian–Volterra kernels produce the singular-value seam and common-channel option response required by the recovery theorem.

13. RESIDUAL KERNEL SELECTION

Source recovery leaves residual pricing in place. The recovered state fixes the canonical coordinates available to a selector; it does not complete the public payoff module or enlarge the tradeable span.

Definition 13.1 (Conditional entropy residual constraint set). Fix a date t and horizon T . Let

$$L_0 \equiv 1, \quad L_1, \dots, L_m \in L^\infty(\mathbb{Q}, \mathcal{K}_t)$$

be the latent loadings of the admitted public payoff constraints, and let $q_1, \dots, q_p \in L^\infty(\mathbb{Q}, \mathcal{K}_t)$ be residual scores whose exposure is to be selected. For a bounded $\mathcal{H}_t$ -measurable target vector $\eta \in L^\infty(\mathbb{Q}, \mathcal{H}_t; \mathbb{R}^p)$, define

$$\mathcal{Z}_t(\eta) := \left\{ Z > 0 : \begin{array}{l} \mathbb{E}^\mathbb{Q}[Z L_j \mid \mathcal{H}_t] = \mathbb{E}^\mathbb{Q}[L_j \mid \mathcal{H}_t], \quad j = 0, \dots, m, \\ \mathbb{E}^\mathbb{Q}[Z q_k \mid \mathcal{H}_t] = \eta_k, \quad k = 1, \dots, p \end{array} \right\}.$$

The conditional entropy of Z is

$$H_t(Z) := \mathbb{E}^\mathbb{Q}[Z \log Z \mid \mathcal{H}_t].$$

Theorem 13.2 (Entropy-minimal residual density). *Assume $\mathcal{Z}_t(\eta)$ is nonempty, satisfies a conditional Slater condition, and has the required conditional exponential moments. Then there is a unique minimizer*

$$Z_t^*(\eta) = \arg \min_{Z \in \mathcal{Z}_t(\eta)} H_t(Z).$$

It has conditional exponential form

$$Z_t^*(\eta) = \exp \left(\alpha_0 + \sum_{j=1}^m \alpha_j L_j + \sum_{k=1}^p \gamma_k q_k \right),$$

where the multipliers are $\mathcal{H}_t$ -measurable and chosen to satisfy the constraints. The selected price

$$\Pi_t^{*,\eta}(X) := \mathbb{E}^{\mathbb{Q}}[Z_t^*(\eta)X \mid \mathcal{H}_t]$$

is linear, positive, public-equivalent on the admitted public constraints, and differentiable in η wherever the conditional covariance matrix of the active constraints is nonsingular.

Proof. Conditional relative entropy is strictly convex on positive densities, and the moment constraints form a convex affine set. The conditional Slater condition gives the supporting Lagrange multipliers and excludes boundary minimizers. Minimizing the conditional Lagrangian pointwise in the conditional problem yields the exponential density displayed above; strict convexity gives uniqueness. Differentiability in η follows from the conditional implicit-function theorem applied to the moment equations, whose Jacobian is the conditional covariance matrix of the active constraints under the tilted density. $\square$

Remark 13.3 (Residual envelopes and canonical selectors). The residual envelope is the admissible price interval. A linear basis selector chooses one point in that interval. The entropy-minimal selector supplies a canonical no-arbitrage choice once public-equivalence constraints and residual target exposures are fixed.

Definition 13.4 (Terminal tradeable payoff span). For a horizon $T > t$, let $\mathcal{S}_{t,T}^{\text{pay}} \subset L^2(\mathbb{Q}, \mathcal{F}_T)$ be the closed subspace of discounted terminal payoffs generated by admissible public self-financing strategies and by the listed or idealized public instruments admitted for trading between t and T . Orthogonal hedgeability statements are taken relative to this terminal payoff span.

Proposition 13.5 (Recovery does not imply tradeable completion). Fix $T > t$. Suppose the option manifold recovers the canonical state $\Xi_t^{\mathbb{Q}}$. If $(\mathcal{S}_{t,T}^{\text{pay}})^{\perp}$ contains a nonzero residual direction correlated with a claim X , and this direction can be chosen to annihilate the admitted public payoff span, then recovering $\Xi_t^{\mathbb{Q}}$ does not make X tradeably replicable from $\mathcal{S}_{t,T}^{\text{pay}}$. Residual kernel selection remains a pricing input.

Proof. Recovery makes the canonical coordinates measurable functions of the option manifold. It does not change the terminal tradeable payoff span $\mathcal{S}_{t,T}^{\text{pay}}$. If the orthogonal residual of X relative to this span is nonzero, the Hilbert projection theorem leaves a strict residual norm after the best tradeable projection. If the residual direction also annihilates the admitted public payoff span, Lemma 6.6 constructs public-equivalent density families that price that direction without changing public prices. A selector is therefore still needed. $\square$

Definition 13.6 (Dynamic residual selector). Let

$$\mathcal{R}_{t,T} = \{R^1, \dots, R^m\} \subset L^\infty(\mathbb{Q}, \mathcal{K}_t)$$

be a finite horizon-specific family of bounded date- t residual directions satisfying

$$\mathbb{E}^{\mathbb{Q}}[R^r \mid \mathcal{H}_t] = 0, \quad r = 1, \dots, m,$$

$$\mathbb{E}^{\mathbb{Q}}[R^r X_a \mid \mathcal{H}_t] = 0 \quad \text{for every admitted public payoff } X_a \text{ and every } r.$$

A dynamic residual selector is an $\mathcal{H}_t \vee \sigma(\Xi_t^{\mathbb{Q}})$ -measurable vector

$$\theta_t = (\theta_{1,t}, \dots, \theta_{m,t})$$

such that

$$1 + \sum_{r=1}^m \theta_{r,t} R^r > 0.$$

The associated density family is horizon-wise; it need not be a dynamically consistent density process unless the martingale consistency conditions in Section 10 are imposed.

Definition 13.7 (Dynamic option-manifold selected residual price). For $X \in L^2(\mathbb{Q}, \mathcal{F}_T)$, define

$$\Pi_t^\theta(X) := \mathbb{E}^\mathbb{Q}[X \mid \mathcal{H}_t] + \sum_{r=1}^m \theta_{r,t} \mathbb{E}^\mathbb{Q}[L_{t,T}^\mathbb{Q}(X)R^r \mid \mathcal{H}_t].$$

This is the linear price induced by the selected density

$$Z_{t,S}^\theta = 1 + \sum_{r=1}^m \theta_{r,t} R^r.$$

Definition 13.8 (Residual price envelope). For a conditionally bounded residual family $\mathfrak{R}_{t,T}$ and admissible radius $\bar{\theta}_t(R) \|R\|_\infty < 1$, define the upper residual envelope by

$$\bar{\Pi}_t(X) := \mathbb{E}^\mathbb{Q}[X \mid \mathcal{H}_t] + \operatorname{ess\,sup}_{R \in \mathfrak{R}_{t,T}, |\eta| \leq \bar{\theta}_t(R)} \eta \mathbb{E}^\mathbb{Q}[L_{t,T}^\mathbb{Q}(X)R \mid \mathcal{H}_t],$$

whenever the conditional essential supremum is finite. The lower envelope is defined by the corresponding essential infimum.

Theorem 13.9 (Selected prices lie in the residual envelope). *Suppose $\mathcal{R}_{t,T} = \{R^1, \dots, R^m\} \subset \mathfrak{R}_{t,T}$ is conditionally bounded in L^∞ and the selector satisfies*

$$|\theta_{r,t}| \leq \bar{\theta}_t(R^r), \quad \left\| \sum_{r=1}^m \theta_{r,t} R^r \right\|_\infty < 1.$$

Then $\Pi_t^\theta(X)$ lies inside the public-equivalent residual price envelope generated by the density families $Z_{t,S}^{\eta,R} = 1 + \eta R$, $|\eta| \leq \bar{\theta}_t(R)$, $R \in \mathfrak{R}_{t,T}$, provided the selected aggregate residual $\sum_r \theta_{r,t} R^r$ is admitted by the envelope family.

Proof. For each R , the density family $Z_{t,S}^{\eta,R} = 1 + \eta R$ is positive and public-equivalent by Lemma 6.6. Its conditional price of X is

$$\mathbb{E}^\mathbb{Q}[X \mid \mathcal{H}_t] + \eta \mathbb{E}^\mathbb{Q}[L_{t,T}^\mathbb{Q}(X)R \mid \mathcal{H}_t].$$

The selected price is the conditional expectation under the single positive density family

$$1 + \sum_r \theta_{r,t} R^r.$$

If the selected aggregate residual belongs to the admissible envelope family, this linear price is one member of the support family over which the envelope takes the essential supremum and infimum. Hence it lies between the lower and upper residual envelopes. $\square$

Definition 13.10 (Residual Gram coordinates). For a finite residual basis $R = (R^1, \dots, R^m)^\top$, define the conditional residual Gram matrix and payoff-residual vector by

$$G_t^R := \mathbb{E}^\mathbb{Q}[RR^\top \mid \mathcal{H}_t], \quad g_t^X := \mathbb{E}^\mathbb{Q}[L_{t,T}^\mathbb{Q}(X)R \mid \mathcal{H}_t].$$

For a selector vector θ_t , the selected residual price can be written

$$\Pi_t^\theta(X) = \mathbb{E}^\mathbb{Q}[X \mid \mathcal{H}_t] + \theta_t^\top g_t^X,$$

provided the selected density is positive and public-equivalent.

Theorem 13.11 (Ellipsoidal residual envelope). *Assume G_t^R is positive definite on the active residual span. For the conditional selector budget*

$$\Theta_t^0(\rho) := \{\theta : \theta^\top G_t^R \theta \leq \rho_t^2\},$$

the unconstrained linear residual support function is

$$\sup_{\theta \in \Theta_t^0(\rho)} \theta^\top g_t^X = \rho_t \sqrt{(g_t^X)^\top (G_t^R)^{-1} g_t^X}.$$

The corresponding lower endpoint is obtained by changing the sign. If the selector budget is intersected with the positivity domain of the selected densities, the same formula remains an upper bound, and equality holds whenever the Cauchy–Schwarz optimizer lies in that positivity

domain. In particular, equality holds for every claim in the active residual span under the small-ball condition

$$\rho_t \sqrt{\text{ess sup } R^\top (G_t^R)^{-1} R} < 1.$$

Therefore, inside the unconstrained ellipsoidal selector class, and inside the constrained class whenever the preceding positivity condition applies,

$$\Pi_t^\theta(X) \in \mathbb{E}^\mathbb{Q}[X \mid \mathcal{H}_t] \pm \rho_t \sqrt{(g_t^X)^\top (G_t^R)^{-1} g_t^X}.$$

Proof. This is Cauchy–Schwarz in the conditional inner product $\langle a, b \rangle_G := a^\top G_t^R b$:

$$\theta^\top g = \left\langle \theta, (G_t^R)^{-1} g \right\rangle_G \leq (\theta^\top G_t^R \theta)^{1/2} (g^\top (G_t^R)^{-1} g)^{1/2}.$$

Equality is achieved in the direction $(G_t^R)^{-1} g^X$, unless restricted by the positivity domain. Intersecting with the positivity domain can only reduce the support function. The small-ball condition implies $1 + \theta^\top R > 0$ almost surely for every $\theta^\top G_t^R \theta \leq \rho_t^2$, again by Cauchy–Schwarz, and hence the unconstrained optimizer is admissible. The lower endpoint follows by the opposite direction. $\square$

Corollary 13.12 (Box residual envelope). *For the unconstrained box selector budget $|\theta_r| \leq \bar{\theta}_r$,*

$$\sup_{|\theta_r| \leq \bar{\theta}_r} \theta^\top g_t^X = \sum_{r=1}^m \bar{\theta}_r |g_{t,r}^X|,$$

and the lower endpoint is obtained by changing the sign. If the box is intersected with the positivity domain of the selected densities, the same expression is an upper bound and equality holds when the coordinatewise sign optimizer is positive.

Proof. The support function of the box $\prod_r [-\bar{\theta}_r, \bar{\theta}_r]$ in direction g_t^X is the sum of the coordinate support functions,

$$\sum_{r=1}^m \sup_{|\theta_r| \leq \bar{\theta}_r} \theta_r g_{t,r}^X = \sum_{r=1}^m \bar{\theta}_r |g_{t,r}^X|.$$

The lower endpoint is the support function in direction $-g_t^X$. Intersecting the box with the positivity domain only removes admissible selectors, so the expression remains an upper bound and is attained whenever the coordinatewise sign optimizer lies in that domain. $\square$

Proposition 13.13 (Adaptedness of selected prices). *If $\Xi_t^\mathbb{Q}$, $\mathcal{R}_{t,T}$, and θ_t are $\mathcal{H}_t \vee \sigma(\Xi_t^\mathbb{Q})$ -measurable and the recovered canonical state is $\mathcal{H}_t$ -measurable through the option manifold, then $\Pi_t^\theta(X)$ is $\mathcal{H}_t$ -measurable.*

Proof. The public price is $\mathcal{H}_t$ -measurable. Under recovery, the canonical coordinates entering θ_t are measurable functions of $\mathcal{M}_t^{\text{opt}}$, hence $\mathcal{H}_t$ -measurable. Conditional pairings with residual directions are conditional expectations given $\mathcal{H}_t$, hence $\mathcal{H}_t$ -measurable. The finite selected sum is therefore $\mathcal{H}_t$ -measurable under the stated finiteness assumption. $\square$

14. NATIVE STRESS GREEKS

Classical option Greeks differentiate prices with respect to public surface coordinates or underlying price variables. Volterra–Perron Greeks differentiate selected prices with respect to recovered quotient coordinates. Once the source quotient is recovered, the primary sensitivity is a cotangent vector on that quotient. Public quote Greeks are adjoint representations of the same object, not separate primitives.

Definition 14.1 (Native stress Greeks). Let $\Pi_t^\theta(X; \Theta)$ be the selected price at state Θ . If the indicated directional derivatives exist, define

$$\begin{aligned} \Delta_\varpi X &:= D\Pi_t^\theta(X; \Theta)[\delta\varpi], & \Delta_\rho X &:= D\Pi_t^\theta(X; \Theta)[\delta\rho], \\ \Delta_{e_i} X &:= D\Pi_t^\theta(X; \Theta)[\delta e_i], & \Delta_{\mathfrak{H}} X &:= D\Pi_t^\theta(X; \Theta)[\delta\mathfrak{H}], \\ \Delta_{G_{ij}} X &:= D\Pi_t^\theta(X; \Theta)[\delta G_{ij}], & \Delta_\theta X &:= D\Pi_t^\theta(X; \Theta)[\delta\theta]. \end{aligned}$$

These are, respectively, Perron-stress, spectral-radius, Perron-eigenvector, roughness, directed-contagion, and residual-selector Greeks. If carrier weights are separate state coordinates, their derivatives define additional carrier-weight Greeks.

Assumption 14.2 (Differentiable selected pricing). In a neighborhood of the recovered state, the map

$$\Theta \mapsto \Pi_t^\theta(X; \Theta)$$

is Fréchet differentiable on a chart transverse to gauge directions. The residual selector is differentiable in the same chart.

Theorem 14.3 (Existence and gauge invariance of native Greeks). *Under Assumption 14.2, native stress Greeks exist as continuous linear functionals on the canonical quotient tangent space. If two chart perturbations differ by a gauge direction, they have the same native Greek.*

Proof. Fréchet differentiability gives a continuous linear derivative on the chart tangent space. The selected price is a function of the gauge class of the state. Gauge transformations leave latent loadings, public option prices, and selected residual prices unchanged. Hence the derivative vanishes on gauge directions and factors through the quotient tangent space. The quotient functional is the native Greek. $\square$

Theorem 14.4 (Recovered-manifold Greek). *Suppose the recovery theorem holds in a C^1 embedded quotient chart, for instance because the quotient source chart is finite dimensional with full column rank or because $D\bar{F}$ has closed complemented range and locally constant split rank. Suppose also that $\Pi_t^\theta(X; \Theta)$ is differentiable in the canonical state. On the recovery neighborhood, the selected price can be written as a function of the public option manifold,*

$$\hat{\Pi}_t^\theta(X; y) := \Pi_t^\theta(X; \bar{F}^{-1}(y)).$$

Its first derivative is

$$D_y \hat{\Pi}_t^\theta(X; y) = D_\Theta \Pi_t^\theta(X; \Theta) \circ D_y \bar{F}^{-1}(y), \quad \Theta = \bar{F}^{-1}(y).$$

Thus a public surface perturbation has a latent-stress decomposition through the recovered Volterra–Perron coordinates.

Proof. The local inverse $\bar{F}^{-1}$ is differentiable by the recovery theorem and the inverse function theorem. The displayed identity is the chain rule applied to the composition $\Pi_t^\theta \circ \bar{F}^{-1}$. $\square$

Definition 14.5 (Counterfactual Volterra–Perron shock). At a recovered state $\Xi_t^\mathbb{Q}$, a counterfactual Volterra–Perron shock is a quotient tangent vector

$$h \in T_{[\Xi_t^\mathbb{Q}]}(\mathfrak{S}_t / \sim_g)$$

that satisfies the linearized no-arbitrage tangent conditions of the public option manifold. Its first-order public surface displacement is

$$\delta_h \mathcal{M}_t^{\text{opt}} := D\mathfrak{P}_t([\Xi_t^\mathbb{Q}])h,$$

and its selected latent-price displacement for claim X is

$$\mathfrak{G}_h(X) := D_\Xi \Pi_t^\theta(X; \Xi_t^\mathbb{Q})[h].$$

Canonical examples are Perron-amplification shocks, carrier-thinning shocks, roughness shocks, directed-route shocks, and residual-selector shocks.

Proposition 14.6 (Gauge-invariant stress calculus). *If the public pricing map and selected latent price are well defined on the quotient state, then both maps*

$$h \mapsto D\mathfrak{P}_t([\Xi_t^\mathbb{Q}])h, \quad h \mapsto D_\Xi \Pi_t^\theta(X; \Xi_t^\mathbb{Q})[h]$$

are independent of the representative chosen for $\Xi_t^\mathbb{Q}$. Before quotienting, both derivatives vanish on pure gauge directions.

Proof. Maps defined on a quotient are constant on gauge orbits. Their derivatives annihilate tangent vectors to those orbits and descend to the quotient tangent space. The two displayed derivatives are exactly the descended derivatives. $\square$

Definition 14.7 (Adjoint quote Greek). Let $F_\tau : U \subset \mathbb{R}^d \rightarrow \mathbb{R}^M$ be a finite-panel observation map with full column-rank Jacobian $J = DF_\tau(u)$. Let $p_X := D_u \Pi_t^\theta(X; u) \in (\mathbb{R}^d)^*$ be a native Greek row vector, and let $W \succ 0$ be the finite-panel quote metric. A quote-space adjoint Greek is any quote covector $q_X \in \mathbb{R}^M$ satisfying

$$J^\top q_X = p_X^\top.$$

The minimum W -norm adjoint Greek, where $\|q\|_W^2 = q^\top W q$, is

$$q_X^{*,W} := W^{-1} J (J^\top W^{-1} J)^{-1} p_X^\top.$$

Theorem 14.8 (Quote perturbation representation). *For any quote perturbation $\delta y \in \text{Ran } J$,*

$$D_y(\Pi_t^\theta \circ F_\tau^{-1})[\delta y] = (q_X^{*,W})^\top \delta y.$$

If $\delta y = J \delta u$, then this equals $p_X \delta u$.

Proof. The formula for $q_X^{*,W}$ is the Lagrange multiplier solution of minimizing $q^\top W q$ subject to $J^\top q = p_X^\top$. Hence $J^\top q_X^{*,W} = p_X^\top$. For tangent perturbations $\delta y = J \delta u$,

$$(q_X^{*,W})^\top \delta y = (q_X^{*,W})^\top J \delta u = (J^\top q_X^{*,W})^\top \delta u = p_X \delta u.$$

This is exactly the native-state derivative of the selected price. $\square$

Corollary 14.9 (Native and public Greeks are adjoint images). *Under the recovery hypotheses of Theorem 8.22, every differentiable selected claim price has a canonical native Greek on the source quotient. On any finite full-rank public panel with quote metric W , its minimum W -norm public quote Greek is the adjoint solution*

$$q_X^{*,W} = W^{-1} J (J^\top W^{-1} J)^{-1} p_X^\top.$$

Thus Black–Scholes Greeks and Greeks computed in public quote coordinates are adjoint images of quotient sensitivities to roughness, directed routes, Perron amplification, carrier absorption, phase transition scores, and residual selection.

Proof. The native Greek is the quotient derivative from Theorem 14.3. On a recovered finite panel, J maps quotient source perturbations to public quote perturbations. The weighted adjoint equation is exactly the representation identity in Theorem 14.8. $\square$

Remark 14.10 (Finite Greek suites). In a finite implementation one may report a named suite of native Greek directions by choosing source-chart tangents for Perron stress, spectral amplification, rough maturity scaling, directed routes, carrier weights, phase scores, and residual-selector shifts. The weighted adjoint equation then turns each declared tangent into a quote-space functional and records any component outside the admitted quote tangent. Such a suite is a certificate-layer object. It measures how the finite option panel represents selected quotient sensitivities, without changing the quotient recovery theorem.

Theorem 14.11 (First-order overidentifying restrictions). *Let $F_\tau : \mathbb{R}^d \rightarrow \mathbb{R}^M$ be a finite-panel source response with $M > d$ and full column-rank Jacobian $J = DF_\tau(u_0)$. Let*

$$\mathcal{T}_{y_0} := \text{Ran } J, \quad \mathcal{N}_{y_0} := \text{Ker } J^\top, \quad y_0 = F_\tau(u_0).$$

If a perturbed quote panel $y_\varepsilon = y_0 + \varepsilon h + o(\varepsilon)$ remains on the model image to first order, then

$$n^\top h = 0 \quad \text{for every } n \in \mathcal{N}_{y_0}.$$

Equivalently, $(I - P_\mathcal{T})h = 0$, where $P_\mathcal{T}$ is Euclidean projection onto $\mathcal{T}_{y_0}$. Thus a finite panel with more quote coordinates than source coordinates imposes $M - d$ first-order normal-space restrictions at a full-rank point.

Proof. If $y_\varepsilon = F_\tau(u_0 + \varepsilon v + o(\varepsilon))$, then

$$h = Jv \in \text{Ran } J.$$

Membership in $\text{Ran } J$ is equivalent to orthogonality to $\text{Ker } J^\top$, and also to $(I - P_\tau)h = 0$. $\square$

Proposition 14.12 (Perron spectral Greek). *Suppose the screened synchronization operator $A^\mathbb{Q}$ has a simple Perron eigenvalue $\rho^\mathbb{Q}$ with normalized left and right eigenvectors $\ell^\mathbb{Q}(e^\mathbb{Q}) = 1$. For a bounded operator perturbation δA ,*

$$\delta \rho^\mathbb{Q} = \ell^\mathbb{Q}(\delta A e^\mathbb{Q}).$$

Consequently, any selected price differentiable in $\rho^\mathbb{Q}$ has spectral sensitivity

$$D\Pi_t^\theta[\delta A] = \partial_\rho \Pi_t^\theta \ell^\mathbb{Q}(\delta A e^\mathbb{Q}) + \text{eigenvector and residual-selector terms.}$$

Proof. Differentiate $Ae = \rho e$ and apply the normalized left eigenfunctional ℓ :

$$\ell(\delta Ae) + \ell(A\delta e) = \delta \rho \ell(e) + \rho \ell(\delta e).$$

Since $\ell A = \rho \ell$ and $\ell(e) = 1$, the $\ell(\delta e)$ terms cancel and $\delta \rho = \ell(\delta Ae)$. Differentiating the selected price through the spectral chart and applying the chain rule gives the second display. $\square$

Proposition 14.13 (Perron eigenvector derivatives). *Let $Ae = \rho e$, $\ell A = \rho \ell$, and $\ell(e) = 1$, with ρ simple and isolated. Let S be the reduced resolvent of $A - \rho I$ on the complement of $e\ell$. For a perturbation δA preserving the chosen normalization,*

$$\delta e = -S(\delta A - \delta \rho I)e, \quad \delta \ell = -\ell(\delta A - \delta \rho I)S.$$

Consequently,

$$|\delta \rho| \leq \|\ell\| \|e\| \|\delta A\|, \quad \|\delta e\| + \|\delta \ell\| \leq C \kappa_P \|\delta A\|,$$

where $\kappa_P := \|S\|$ is the reduced-resolvent norm and C depends only on the local normalization. On a finite-dimensional Perron stratum with uniformly bounded eigenprojection conditioning, $\kappa_P \leq C'/\gamma_P$; without such conditioning, κ_P is the correct stability constant.

Proof. Differentiating $Ae = \rho e$ gives

$$(A - \rho I)\delta e = -(\delta A - \delta \rho I)e.$$

The normalization fixes the component of δe in the Perron direction. Inverting $A - \rho I$ on the complement of $e\ell$ gives the displayed formula for δe . The left-eigenvector identity follows in the same way from differentiating $\ell A = \rho \ell$ and imposing $\delta \ell(e) + \ell(\delta e) = 0$. The norm bound is controlled by the reduced-resolvent norm $\kappa_P = \|S\|$; the gap estimate is a conditioned finite-dimensional corollary, not the primitive constant. $\square$

15. SYNCHRONIZATION RELATIVE VALUE AND PUBLIC-NEUTRAL HEDGING

Let I be an index with weights ω_i . The index variance decomposition separates single-name variance from synchronized covariance. In recovered coordinates, the same split becomes a selected pricing identity.

Definition 15.1 (Recoverability, tradeability, and hedgeability). Fix $t < T$. The state is recoverable when the public response is locally invertible on the Volterra–Perron source quotient. A claim is tradeable if its payoff lies in the admitted listed or idealized public instrument span at horizon T . A claim is hedgeable if its centered payoff lies in the closed gain space generated by admissible self-financing strategies between t and T .

Theorem 15.2 (Recovery does not imply tradeable hedgeability). *Assume the option manifold recovers the canonical state $\Xi_t^\mathbb{Q}$. Then the following hold.*

- (1) *recovery of $\Xi_t^\mathbb{Q}$ does not imply that a Perron stress claim or residual-stress claim belongs to the tradeable public span;*
- (2) *residual-kernel selection can price latent-stress claims while leaving them unspanned by tradeable instruments;*

- (3) a public-neutral book can have zero exposure to selected public surface directions and nonzero exposure to a recovered synchronization direction, but this exposure is hedgeable only under an additional traded-spanning assumption.

Consequently,

$$\text{recoverability} \neq \text{tradeability} \neq \text{hedgeability}.$$

The first is a quotient-inversion property, the second is an instrument-spanning property, and the third is a dynamic gain-space property.

Proof. The recovery theorem identifies a canonical source coordinate from public option prices. It does not enlarge the closed terminal payoff span generated by tradeable instruments, nor does it alter the dynamic gain space of self-financing strategies. If a latent-stress payoff has a nonzero orthogonal component relative to the terminal tradeable span, the Hilbert projection theorem leaves a strictly positive hedge residual after the best tradeable projection. Residual selectors change prices through public-annihilating directions, but public equivalence only preserves admitted public prices; it does not construct a gain representation. Recoverability is therefore quotient inversion, tradeability is payoff-spanning, and hedgeability is membership in the dynamic gain space. $\square$

Remark 15.3 (Carrier resolution). In finite option manifolds, tradeability also depends on the resolution of the admitted carrier span. A coarse underlier carrier library, an index/sector carrier library, a single-name option-coordinate carrier library, and a regularized coordinate book need not span the same synchronization directions. Carrier absorption is a finite-span property of the chosen public instrument system, not a consequence of latent-state recoverability alone.

Definition 15.4 (Carrier-replicated value). For an index option or variance claim X_I , let X_{car} be the carrier-replicated single-name claim formed from the admitted carrier weights and matching maturity coordinates. Let $\varpi_{t,T}^{\mathbb{Q}}$ be the horizon-lifted synchronized stress direction. In the scalar case define

$$X_{\text{sync}} := \frac{\mathbb{E}^{\mathbb{Q}}[(X_I - X_{\text{car}})\varpi_{t,T}^{\mathbb{Q}} \mid \mathcal{H}_t]}{\mathbb{E}^{\mathbb{Q}}[(\varpi_{t,T}^{\mathbb{Q}})^2 \mid \mathcal{H}_t]} \varpi_{t,T}^{\mathbb{Q}},$$

with the convention that the coefficient is zero if the conditional denominator vanishes. For a multidimensional stress subspace, replace the scalar ratio by the corresponding conditional Gram inverse. The tradeable basis error is

$$B_t(X_I) := X_I - X_{\text{car}} - X_{\text{sync}},$$

so the synchronized component is canonical once the carrier book and stress subspace are fixed.

Definition 15.5 (Synchronization premium and residual charge). The synchronization premium is

$$\text{SynPrem}_t(X_I) := \mathbb{E}^{\mathbb{Q}}[X_{\text{sync}} \mid \mathcal{H}_t].$$

The residual compression charge selected by θ is

$$\text{ResCharge}_t^{\theta}(X_I) := \Pi_t^{\theta}(X_{\text{sync}}) - \mathbb{E}^{\mathbb{Q}}[X_{\text{sync}} \mid \mathcal{H}_t].$$

Remark 15.6 (Public shadow versus pure residual synchronization). The stress direction $\varpi_{t,T}^{\mathbb{Q}}$ is centered by convention, but it need not be orthogonal to $\mathcal{H}_t$. Its public compression may therefore carry a nonzero synchronization premium. If one chooses instead the purely public-residual stress direction

$$\varpi_{t,T}^{\mathbb{Q}} - \mathbb{E}^{\mathbb{Q}}[\varpi_{t,T}^{\mathbb{Q}} \mid \mathcal{H}_t],$$

then the public synchronization premium of that residual component vanishes and the selected price contribution appears through the residual charge. The identity below allows both conventions; it separates the public shadow of synchronized stress from the residual public-equivalent pricing selection.

Theorem 15.7 (No-arbitrage synchronization relative-value identity). *Suppose the selected price on the span of $\{X_I, X_{\text{car}}, X_{\text{sync}}, B_t(X_I)\}$ is induced by a fixed positive public-equivalent density family $Z_{t,\cdot}^\theta$. Then the selected valuation is a no-arbitrage extension of the public option manifold on this span. For every square-integrable index claim X_I with decomposition*

$$X_I = X_{\text{car}} + X_{\text{sync}} + B_t(X_I),$$

the selected price satisfies

$$\Pi_t^\theta(X_I) - \Pi_t^\theta(X_{\text{car}}) = \text{SynPrem}_t(X_I) + \text{ResCharge}_t^\theta(X_I) + \Pi_t^\theta(B_t(X_I)).$$

Proof. Because $Z_{t,\cdot}^\theta$ is positive and public-equivalent, conditional expectation under the horizon component $Z_{t,T}^\theta$ is a positive linear extension of the public price system and preserves all date- t public option prices. The extension is therefore arbitrage-free on the finite span under consideration. Applying this linear functional to $X_I = X_{\text{car}} + X_{\text{sync}} + B_t(X_I)$ gives the decomposition of $\Pi_t^\theta(X_I)$. Subtracting $\Pi_t^\theta(X_{\text{car}})$ leaves the selected price of X_{sync} and the basis term. Splitting $\Pi_t^\theta(X_{\text{sync}})$ into its public conditional expectation and selected residual increment gives SynPrem_t and $\text{ResCharge}_t^\theta$. $\square$

Definition 15.8 (Model public-neutral book). Let $H_t^{\text{pub}} \subseteq L^2(\mathbb{Q}, \mathcal{H}_t)$ be the finite public surface span used for neutrality. A listed or idealized book $Y_t \in L^2(\mathbb{Q})$ is model public-neutral if

$$C_{H_t^{\text{pub}}}^\mathbb{Q} Y_t = 0.$$

It is synchronization-exposed if

$$\langle Y_t, \varpi_t^\mathbb{Q} \rangle_{L^2(\mathbb{Q})} \neq 0$$

after projecting $\varpi_t^\mathbb{Q}$ into the claim horizon.

Theorem 15.9 (Public-neutral residual hedge). *Let X_I be an index claim and let*

$$\mathcal{B}_t = Y_0 + \mathcal{V}_t$$

be a nonempty closed affine class of books satisfying the public-neutrality constraint, with $\mathcal{V}_t$ a closed linear subspace of $L^2(\mathbb{Q})$. Then the minimum-residual book

$$Y_t^\star = \arg \min_{Y \in \mathcal{B}_t} \|X_{\text{sync}} - Y\|_{2,\mathbb{Q}}$$

exists and is unique. The residual

$$X_{\text{sync}} - Y_t^\star$$

is orthogonal to feasible public-neutral perturbations of the book. If additional exposure caps make $\mathcal{B}_t$ closed and convex rather than affine, the projection still exists and is unique, and the orthogonality condition is replaced by the usual variational inequality.

Proof. The Hilbert projection theorem gives a unique metric projection of X_{sync} onto the nonempty closed affine subspace $\mathcal{B}_t$. For an affine feasible class, first-order optimality is exactly

$$X_{\text{sync}} - Y_t^\star \perp \mathcal{V}_t.$$

For a closed convex feasible class, the same projection theorem gives existence and uniqueness, and first-order optimality becomes the variational inequality

$$\langle X_{\text{sync}} - Y_t^\star, Y - Y_t^\star \rangle_{L^2(\mathbb{Q})} \leq 0, \quad Y \in \mathcal{B}_t.$$

$\square$

16. EQUIVALENT-MARTINGALE-MEASURE SPECIALIZATION

The preceding theory works with adapted arbitrage-free option manifolds. This section records one equivalent-martingale-measure specialization in a continuous-time risk-neutral stock-volatility system.

Assumption 16.1 (Reduced risk-neutral stock-volatility system). For $i = 1, \dots, N$, under $\mathbb{Q}$,

$$\frac{dS_i(t)}{S_i(t)} = r_t dt + \sqrt{V_i(t)} dW_i^{\mathbb{Q}}(t),$$

where $V_i(t) = \chi_i(Y_i(t))$ for a locally Lipschitz map $\chi_i : \mathbb{R} \rightarrow [0, \infty)$, such as an exponential or softplus transform. The stress factor vector satisfies the Volterra equation

$$Y_i(t) = y_i^0(t) + \sum_{a=1}^m \int_0^t K_{ia}^{\mathbb{Q}}(t-s) dB_a^{\mathbb{Q}}(s) + \sum_{j \neq i} \int_0^t G_{ij}^{\mathbb{Q}}(t-s) \phi_j(Y_j(s)) ds.$$

The kernels are square-integrable on compact horizons, the functions ϕ_j are Lipschitz with linear growth, and Y admits a weak solution with $\mathbb{E}^{\mathbb{Q}} \int_0^T V_i(t) dt < \infty$. The Brownian vectors $(W^{\mathbb{Q}}, B^{\mathbb{Q}})$ have a specified joint covariance; in the reduced examples below this covariance is fixed as part of the source chart.

Theorem 16.2 (Risk-neutral option manifold under an EMM). *Under Assumption 16.1, discounted asset prices are local $\mathbb{Q}$ -martingales. If they are true martingales on $[0, T]$, then European option prices*

$$C_t(a) = \mathbb{E}^{\mathbb{Q}} \left[e^{-\int_t^{T(a)} r_s ds} X_a \mid \mathcal{H}_t \right]$$

form an arbitrage-free public option manifold. The directed Volterra kernel $G^{\mathbb{Q}}$ and the screened Perron operator derived from it define a Volterra–Perron state in the sense of Definition 3.1.

Proof. The discounted stock dynamics have zero drift under $\mathbb{Q}$, so discounted prices are local martingales. Under the true-martingale assumption, conditional-expectation pricing gives an arbitrage-free price system on the admitted European payoff span. Square-integrability of the kernels and the weak solution assumptions place the directed Volterra operator in the Hilbert space setting used throughout the analysis. Screening and Perron projection then produce the spectral triple whenever the positive screened operator has a simple Perron eigenvalue. $\square$

Assumption 16.3 (Reduced normal-forward Volterra system). For each underlier i and maturity T , let $F_i(s, T)$, $0 \leq s \leq T$, be the T -forward price under the T -forward measure $\mathbb{Q}^T$. On the pricing window,

$$dF_i(s, T) = \sum_{a=1}^m \sigma_{ia}^N(s, T, Y_s) dW_a^T(s), \quad 0 \leq s \leq T,$$

where the normal-forward volatility coefficients are locally square-integrable and adapted. Their latent stress state is generated by the same directed Volterra source family as in Assumption 16.1, possibly after replacing the positive variance transform by a normal-volatility transform. The money-market account is strictly positive and the discount curve is public. In the finite benchmark, the normal covariance is

$$\int_t^T \Sigma^N(s, T, Y_s) \Sigma^N(s, T, Y_s)^\top ds = \Gamma_t(T - t; U_t)$$

after conditioning on $\mathcal{H}_t$ and freezing the recovered state at date t , or exactly when the displayed integrated covariance is $\mathcal{H}_t$ -measurable.

Proposition 16.4 (Normal-forward option manifold under an EMM). *Under Assumption 16.3, a T -settled European payoff X_T has the arbitrage-free representation*

$$C_t(X_T) = P(t, T) \mathbb{E}^{\mathbb{Q}^T} [X_T \mid \mathcal{H}_t].$$

If a one-period conditional Gaussian approximation is used in the recovered state chart, the local price coordinates are Bachelier coordinates with normal total variance

$$v_{u,t}^N(T) = \text{Var}^{\mathbb{Q}^T}(F_u(T, T) - F_u(t, T) \mid \mathcal{H}_t),$$

and the response ranks are the chart-invariant ranks of Theorem 4.9.

Proof. The forward dynamics have zero drift under $\mathbb{Q}^T$, so forward prices are local martingales in the T -numeraire. The change-of-numeraire formula gives $C_t = P(t, T)\mathbb{E}^{\mathbb{Q}^T}[X_T \mid \mathcal{H}_t]$ for admitted T -settled claims under the usual true-martingale integrability condition. Conditional Gaussian one-period coordinates reduce these prices to Definition 4.7 in the normal chart. The rank statement follows because the Bachelier chart has strictly positive variance vega away from zero normal variance. $\square$

17. REDUCED EXAMPLES

17.1. Two-name Gaussian Volterra recovery. Consider two names with Gaussian variance-stress factors

$$V_1(t) = \int_0^t K_1(t-s) dB_1(s), \quad V_2(t) = \int_0^t K_2(t-s) dB_2(s) + \gamma \int_0^t G_{21}(t-s) dB_1(s).$$

The positive variance entering the stock model may be a transform of these factors as in Assumption 16.1. The cross-kernel coefficient γ is latent. A single marginal surface for name 2 contains the law of V_2 after public compression, but it does not identify whether the additional variance comes from an idiosyncratic route, a directed route from name 1, or a public-equivalent residual tilt. This is the non-identification theorem in its smallest form.

Now impose a Gaussian–Volterra source family in which the singular profile of G_{21} is known up to the scalar coefficient and the common-channel sign convention is fixed. Let

$$R_{12}(\tau) := \text{Cov}^{\mathbb{Q}}\left(\int_t^{t+\tau} V_1(s) ds, \int_t^{t+\tau} V_2(s) ds \mid \mathcal{H}_t\right)$$

be the cross-maturity response extracted from the index/single-name quadratic option block. In the source family,

$$R_{12}(\tau) = \gamma q_{21}(\tau) + r_{\text{low}}(\tau),$$

where q_{21} is the leading Volterra singular-profile response and r_{low} is lower order uniformly on the maturity interval. If q_{21} is not identically zero, the map $\gamma \mapsto R_{12}$ has nonzero derivative in the quotient chart. An open maturity interval, or a finite analytic maturity panel as in Theorem 8.16, therefore recovers γ up to the fixed gauge convention. Public compression alone fails; source-family maturity structure supplies the required rank.

17.2. Index plus carrier manifold. Let an index have weights ω_i . The implied index total variance can be written in the schematic form

$$v_I^{\text{imp}}(\tau) = \sum_i \omega_i^2 v_i^{\text{imp}}(\tau) + D_\omega(v^{\text{imp}}(\tau)) \rho_{\text{sync}}^{\text{imp}}(\tau),$$

where D_ω is the public synchronization denominator. If the carrier surface span is nondegenerate, with Gram matrix $\Gamma_{\text{car}} \succeq \kappa_{\text{car}} I$, the observed family

$$\tau \mapsto (v_I^{\text{imp}}(\tau), v_i^{\text{imp}}(\tau))$$

over an open maturity interval identifies the implied synchronization curve $\rho_{\text{sync}}^{\text{imp}}(\tau)$. A Volterra–Perron source family then represents that curve as the public image of $(\rho^{\mathbb{Q}}, e^{\mathbb{Q}}, \ell^{\mathbb{Q}}, P^{\mathbb{Q}})$. The carrier Gram lower bound is what prevents the single-name side from becoming an ill-conditioned explanation of the index response.

The remaining price difference between the index claim and its carrier-replicated single-name claim decomposes as

$$\Pi_t^\theta(X_I) - \Pi_t^\theta(X_{\text{car}}) = \text{SynPrem}_t(X_I) + \text{ResCharge}_t^\theta(X_I) + \Pi_t^\theta(B_t(X_I)).$$

The first term is the public synchronized-stress premium read from the option manifold. The second is the selected residual compression charge attached to unspanned latent stress. The third is the tradeable basis error left by the carrier span. Thus the same manifold that recovers the canonical synchronization state also tells the model where recovery stops and pricing selection begins.

17.3. Three-name carrier-rank example. For $N = 3$, write the active unordered edge vector as

$$V_{\mathcal{E}}(\tau) = (V_{12,t}^{\mathbb{Q}}(\tau), V_{13,t}^{\mathbb{Q}}(\tau), V_{23,t}^{\mathbb{Q}}(\tau))^{\top}.$$

One broad index with weights $w = (w_1, w_2, w_3)$ encodes only

$$B_I(\tau) = 2w_1w_2V_{12,t}^{\mathbb{Q}}(\tau) + 2w_1w_3V_{13,t}^{\mathbb{Q}}(\tau) + 2w_2w_3V_{23,t}^{\mathbb{Q}}(\tau),$$

which is one equation for three cross routes. Adding three pair carriers with weights

$$b^{12} = (1, 1, 0), \quad b^{13} = (1, 0, 1), \quad b^{23} = (0, 1, 1)$$

gives a basket design matrix proportional to the identity on the three active pairs. The three cross routes are then identified at each maturity. Carrier rank is the algebraic condition that separates synchronization routes rather than only their weighted aggregate.

17.4. Directed-route seam example. With two names, suppose the ordered routes enter a public cross response as

$$R_{12}(\tau) = \gamma_{1 \leftarrow 2}f(\tau) + \gamma_{2 \leftarrow 1}f(\tau).$$

Then only the sum $\gamma_{1 \leftarrow 2} + \gamma_{2 \leftarrow 1}$ is identified. If instead

$$R_{12}(\tau) = \gamma_{1 \leftarrow 2}f_{12}(\tau) + \gamma_{2 \leftarrow 1}f_{21}(\tau),$$

and f_{12} and f_{21} are linearly independent analytic maturity functions, an open interval or a generic two-point panel separates the two directions. This is the smallest illustration of the direction-resolving block in Theorem 7.15.

17.5. Residual-envelope calculation. Let R^1, R^2 be conditionally orthonormal residual directions, so $G_t^R = I_2$, and suppose

$$g_t^X = (g_1, g_2)^{\top}.$$

Under the ellipsoidal residual budget $\theta_1^2 + \theta_2^2 \leq \rho^2$, Theorem 13.11 gives

$$\left[\mathbb{E}^{\mathbb{Q}}[X \mid \mathcal{H}_t] - \rho\sqrt{g_1^2 + g_2^2}, \mathbb{E}^{\mathbb{Q}}[X \mid \mathcal{H}_t] + \rho\sqrt{g_1^2 + g_2^2} \right].$$

Under the box budget $|\theta_i| \leq \bar{\theta}_i$, Corollary 13.12 gives

$$\left[\mathbb{E}^{\mathbb{Q}}[X \mid \mathcal{H}_t] - \bar{\theta}_1|g_1| - \bar{\theta}_2|g_2|, \mathbb{E}^{\mathbb{Q}}[X \mid \mathcal{H}_t] + \bar{\theta}_1|g_1| + \bar{\theta}_2|g_2| \right].$$

The residual compression charge is the support function of the admissible selector set evaluated at the payoff-residual vector. The calculation is a pricing-fiber computation, not a recovery claim.

18. CONCLUSION

The option manifold identifies a quotient, not the full latent risk-neutral state. The governing map is

$$\text{public option prices} \rightsquigarrow \mathfrak{X}_t^{\text{pub}}, \quad \text{state recovery} = \mathfrak{S}_t / \sim_g \hookrightarrow \mathfrak{X}_t^{\text{pub}}.$$

Stable recovery is positivity of the quotient observability modulus. Residual prices collapse to singletons exactly at public payoff-module completion. Dynamic admissibility adds convex-order and recovered martingale drift restrictions. Recovery, completion, and hedgeability remain separate.

The same quotient records the failure modes. Marginal surfaces leave synchronization unresolved. Finite basket panels leave ellipsope uncertainty unless their design spans the active covariance routes. Carrier Gram collapse makes public-neutral dispersion unstable. Perron reduced-resolvent blow-up, including gap collapse, makes the dominant stress mode non-smooth.

Maturity-rank collapse prevents Volterra source identification. Public payoff incompleteness leaves residual pricing fibers. Convex-order or martingale-drift failure rules out a dynamically admissible risk-neutral model. These are local non-identification mechanisms in the option manifold, not estimation artifacts.

Once recovered and paired with an admissible residual selector, the state supports selected prices beyond the public surface. Native Greeks measure sensitivity to Perron stress, spectral radius, carrier direction, roughness, directed contagion, and residual selection. The synchronization relative-value identity expresses the gap between index options and carrier-replicated single-name options as synchronized-stress premium, residual compression charge, and basis error. The finite Bachelier–Volterra–Perron benchmark gives a minimal pricing equation for this geometry: a recovered finite state, a closed option formula, native state Greeks, and a computable public hedge residual. The benchmark is the finite chart in which the quotient recovery theory becomes a tractable no-arbitrage pricing object, rather than a full option-manifold model.

REFERENCES

- [1] 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.
- [2] 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.
- [3] J. Gatheral, T. Jaisson, and M. Rosenbaum, *Volatility is rough*, Quantitative Finance **18** (2018), no. 6, 933–949, doi:10.1080/14697688.2017.1393551.
- [4] O. El Euch and M. Rosenbaum, *The characteristic function of rough Heston models*, Mathematical Finance **29** (2019), no. 1, 3–38, doi:10.1111/mafi.12173.
- [5] E. Abi Jaber, M. Larsson, and S. Pulido, *Affine Volterra processes*, Annals of Applied Probability **29** (2019), no. 5, 3155–3200, doi:10.1214/19-AAP1477.
- [6] E. Abi Jaber and O. El Euch, *Multifactor approximation of rough volatility models*, SIAM Journal on Financial Mathematics **10** (2019), no. 2, 309–349, doi:10.1137/18M1170236.
- [7] B. Horvath, A. Muguruza, and M. Tomas, *Deep learning volatility: a deep neural network perspective on pricing and calibration in (rough) volatility models*, Quantitative Finance **21** (2021), no. 1, 11–27, doi:10.1080/14697688.2020.1817974.
- [8] 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.
- [9] 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.
- [10] M. Schweizer, *Mean-variance hedging for general claims*, The Annals of Applied Probability **2** (1992), no. 1, 171–179, doi:10.1214/aoap/1177005776.
- [11] J. Vidal Llauradó (2026). *Risk-Neutral Compression and the Geometry of Latent Market Stress*. SSRN. <https://doi.org/10.2139/ssrn.6680938>.
- [12] J. Vidal Llauradó (2026). *Perron Stress Pricing under Public Compression: Residual Price Intervals, Synchronized Stress, and Visible Hedging*. SSRN. <https://doi.org/10.2139/ssrn.6700598>.
- [13] D. Williams, *Probability with Martingales*, Cambridge Mathematical Textbooks, Cambridge University Press, 1991, doi:10.1017/CBO9780511813658.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] D. T. Breeden and R. H. Litzenberger, *Prices of state-contingent claims implicit in option prices*, Journal of Business **51** (1978), no. 4, 621–651, doi:10.1086/296025.
- [19] L. Bachelier, *Théorie de la spéculation*, Annales scientifiques de l’École Normale Supérieure **17** (1900), 21–86, doi:10.24033/asens.476.
- [20] 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.
- [21] 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.
- [22] A. Berman and R. J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, Classics in Applied Mathematics, vol. 9, SIAM, 1994, doi:10.1137/1.9781611971262.

- [23] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed., Cambridge University Press, 2012, doi:10.1017/CBO9781139020411.
- [24] 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>.
- [25] J. Vidal Llauradó (2026). *Dynamic Observability of Latent Contagion: Sequential Revelation and Observable Screening*. SSRN. <https://doi.org/10.2139/ssrn.6539199>.
- [26] J. Vidal Llauradó (2026). *Gaussian Seams in Risk-Neutral Compression: A Volterra Source-Family Companion*. Companion A.

Email address: `vidal.llaurado@gmail.com`