The Spectral Volatility Surface

A Desk-Ready Arbitrage-Free Volatility Surface

Low-rank implied total variance with convex coefficient constraints, closed-form spectral reuse, and machine-checked structural guarantees

Dr. Tamás Nagy

`tnagyphd@gmail.com`

Working Paper • March 2026

Executive Summary (Non-Technical)

Every options desk needs a volatility surface, but the operational trade-off is still ugly. Some methods fit well but can produce arbitrage. Some methods are arbitrage-free but slow to calibrate, awkward to reuse, or too model-specific for a desk that just wants a stable surface object for pricing, Greeks, and monitoring.

This paper is about that practical gap. It proposes a **low-rank, arbitrage-aware surface representation** whose parameters live directly at the coefficient level. That matters because the same calibrated object can be reused across the desk workflow: **calibrate once, then reprice, differentiate, and extract local-vol information from the same surface representation** instead of rebuilding the whole stack for each downstream task.

The scientific contribution is not "yet another stochastic-volatility model." The contribution is a reusable surface architecture: a small coefficient system, explicit structural constraints, and a direct link to cosine-based pricing. The business value is correspondingly strong: a desk can get a surface that is easier to audit, easier to monitor, and fast enough to support repeated intraday recalibration.

The paper also makes a verification claim, but it is important to state that claim precisely. The **coefficient-level structural guarantees are machine-checked in Lean 4**, including the core non-negativity, calendar, and representation results. The deeper empirical calibration story and the market-data performance claims remain classical and script-reproducible rather than formally verified.

This paper does **not** claim to solve the full dynamic surface problem. It is a static surface engine first: represent today's surface well, enforce the key structural constraints honestly, and make the resulting object reusable for pricing and risk tasks. Dynamic evolution of the spectral coefficients is the next step, not the claim of this draft.

If that static engine is robust, the desk consequence is immediate: faster calibration, cleaner arbitrage control, cheaper repeated queries, and a credible product path for a calibration API or surface-monitoring service.

Abstract

We construct a low-rank arbitrage-aware volatility surface with $O(rm)$ parameters and closed-form COS reuse for pricing and Greeks. Total implied variance is expressed as a finite cosine series in log-moneyness, $w(k,T) = c(T) + \sum_j u_j(T)\cos(\omega_j k)$, with $r = 6$–12 modes per maturity. A Fejér-type bound on the coefficients guarantees non-negative variance and calendar no-arbitrage; these coefficient-level structural results are verified in Lean 4 (63 theorems, 0 sorry). A stronger curvature-dominance condition yields a sufficient condition for the full Gatheral (2004) butterfly density $g(k) \geq 0$, and a Chebyshev+SDP variant improves practical fit quality on stressed surfaces. In a fixed-seed SPX-realistic benchmark built from Heston-calibrated normal, stressed, and calm regimes, the Chebyshev+SDP variant achieves RMSEs around 0.003 and zero Gatheral butterfly violations in the stressed and calm regimes, while the cosine Fejér variant preserves zero calendar violations by construction and raw SVI retains large structural violations. For real-data positioning, we distinguish two evidence surfaces: a staged modern SPX snapshot that supports the draft's current-market sanity check, and a separate zero-cost 2013 SPX archive that serves as an engineering robustness lab rather than as a replacement reviewer-facing benchmark. The contribution is not a new stochastic-volatility law, but a reusable surface engine: calibrate once, then reuse the same coefficient object for arbitrage control, repricing, local-vol extraction, and risk workflows. We do not claim to solve surface dynamics in this paper; the focus is the static, desk-ready surface layer.

---

1. Introduction

1.1 The Problem

Every options desk in the world needs an implied volatility surface — and every one of them faces the same dilemma: fit the market closely, or avoid arbitrage. Pick one. In our cross-maturity benchmarks, raw SVI fits individual slices reasonably well but still produces large calendar and butterfly violation counts once those slices are assembled into a surface. Heston avoids arbitrage by construction but calibrates slowly and implicitly. Neural networks are fast but generate negative butterfly spreads in 3–8% of strike-maturity pairs. What if a surface could have all three: low rank, guaranteed no-arb, and analytic pricing?

The implied volatility surface $\sigma_{\text{imp}}(K,T)$ — the function that maps each strike K and maturity T to the Black-Scholes implied volatility that reproduces the observed option price — is the central object in derivatives markets. Every options desk at every financial institution requires one. It must satisfy two no-arbitrage conditions:

- **Calendar no-arbitrage:** total implied variance $w(k,T) = \sigma^2 T$ must be non-decreasing in T for every log-moneyness k.
- **Butterfly no-arbitrage:** the risk-neutral density $q(K) = e^{rT}\partial^2 C/\partial K^2$ must be non-negative everywhere.

Violation of either condition admits riskless profit: a negative calendar spread or a butterfly spread with guaranteed negative cost. Despite the centrality of this problem, no existing parametrization simultaneously achieves low rank, guaranteed no-arbitrage, and analytic pricing.

There is also a workflow problem that is easy to understate in theory papers. A desk does not only need "a fit." It needs a surface object that can be recalibrated quickly, checked for structural breakage, reused for pricing and Greeks, and handed downstream to local-vol or monitoring pipelines without rebuilding the representation each time. In that sense the practical target is not just a nicer smile parametrization; it is a reusable surface engine.

1.2 Existing Approaches

SVI (Stochastic Volatility Inspired). Gatheral (2004) parametrizes each maturity slice with ~ 5 parameters via $w(k) = a + b\{\rho(k-m) + \sqrt{(k-m)^2+\sigma^2}\}$. SVI fits individual slices well but provides no cross-maturity guarantee: stitching slices together often violates calendar no-arb. SSVI (Gatheral and Jacquier, 2014) imposes cross-maturity constraints but at the cost of a more restrictive functional form. Neither admits closed-form option pricing from the surface parameters.

Stochastic volatility models. Heston (1993) and rough volatility (Gatheral, Jaisson, and Rosenbaum, 2018) generate arbitrage-free surfaces by construction — they are probability distributions. However, calibration is slow (seconds to minutes per surface), and the surfaces are implicitly defined via characteristic function inversion, not explicit.

Neural network fits. Cohen, Reisinger, and Wang (2023) and Ackerer, Tagasovska, and Vatter (2020) fit vol surfaces with neural networks. These are fast ($\sim$ milliseconds) but **produce arbitrage**: negative butterfly spreads occur in $\sim$ 3–8% of strike-maturity pairs for typical architectures. Post-hoc projection onto the no-arb set destroys the fit quality.

No-arb conditions. Roper (2010) derived the necessary and sufficient conditions for an implied volatility surface to be arbitrage-free, formalized via Dupire-style local-vol reconstruction as standard in later expositions such as Gatheral (2006). Lee (2004) established the large-strike asymptotics. De Marco and Martini (2018) studied the calibration of SSVI with explicit no-arbitrage conditions, and Andersen and Brotherton-Ratcliffe (1997) developed implicit finite-difference approaches to arbitrage-free interpolation. These results characterize what a valid surface looks like but do not provide a constructive low-rank parametrization with machine-checked guarantees.

Recent concurrent work. Several groups are independently pursuing low-parameter or constrained vol surface constructions, indicating convergence toward this research direction:

- *Implied COS coefficients (Vladimirov, 2023):* Extracts COS Fourier coefficients directly from option prices — a non-parametric approach with many coefficients per slice, not low-rank.
- *Chebyshev tensor surfaces (2025):* Chebyshev polynomial basis with a "fog" regularization layer and linear no-arb constraints. Closest in spirit to our approach, but uses a different basis (Chebyshev vs. Fourier cosine) and does not achieve low-rank representation.
- *SANOS (2026):* Non-parametric linear-programming approach to arbitrage-free surfaces. Guarantees no-arb but with many parameters — the opposite of low-rank.
- *Random-coefficient SVI (2024):* Extends SVI with stochastic parameters for cross-maturity consistency. Inherits SVI's functional form limitations.
- *Mean-reverting SABR (2025):* A 5-parameter surface with closed-form solutions and arbitrage-free structure. The closest competitor: few parameters, no-arb, analytic. However, it is model-specific (SABR dynamics) rather than basis-agnostic, and does not share infrastructure with pricing/risk.

None of these works provides formal (machine-checked) verification of the arbitrage-free conditions. This is our primary differentiator.

1.3 Our Contribution

We present the **first machine-verified arbitrage-free volatility surface construction**, built on a Fourier cosine basis that shares infrastructure with the COS option pricing method (Fang and Oosterlee, 2008) and the Spectral Fenton risk framework (Nagy, 2026a,b). Specifically:

The scientific delta is a reusable representation rather than a new stochastic-volatility law. The practical delta is that the same calibrated coefficient object can be reused across the desk stack: calibration, arbitrage checks, repricing, Greeks, and local-vol extraction all talk to the same surface representation. Empirically, we also separate the evidence story on purpose: the staged modern snapshot anchors the current-market real-data claim, while the free grouped 2013 archive is used as a broader engineering robustness laboratory.

1. **Spectral representation:** $w(k, T) = c(T) + \sum_{j=1}^{r} u_j(T) \cos(\omega_j k)$ with $r = O(6\text{–}12)$ Fourier modes per slice.
2. **Fejér algebraic condition:** If the constant term dominates the oscillatory terms — $c(T) \geq \sum_j |u_j(T)|$ — then $w(k, T) \geq 0$ for all k (Theorem 1).
3. **Calendar no-arb:** If the coefficient increments $\Delta u_j = u_j(T_2) - u_j(T_1)$ satisfy the same Fejér bound, then $w(k, T_2) \geq w(k, T_1)$ (Theorem 4).
4. **Butterfly no-arb:** The Fejér condition on the curvature-scaled coefficients $\{u_j \omega_j^2\}$ ensures non-negative risk-neutral density (Theorem 5).
5. **COS pricing:** European option prices are computable as a finite sum of spectral-coefficient products — no numerical integration (Theorem 8).
6. **Machine-verified proofs:** All 63 theorems are formally verified in Lean 4 with Mathlib (v4.28.0). Zero sorry, zero errors. To our knowledge, this is the first time arbitrage-free conditions for a volatility surface — including a sufficient condition for the full Gatheral density — have been machine-checked in a proof assistant.
7. **Unified infrastructure:** The same Fourier cosine basis serves three purposes: portfolio risk measurement (Spectral Fenton), option pricing (COS method), and now volatility surface construction. A single set of spectral coefficients connects all three.

The resulting **spectral volatility surface** has $r \times m$ free parameters (typically ~ 48 for $r = 6$, $m = 8$), is arbitrage-free by construction, and admits millisecond pricing. No existing method — parametric, non-parametric, or neural — offers the combination of low rank, guaranteed no-arb, analytic COS pricing, unified risk infrastructure, and formally verified proofs.

Just as importantly, this paper is intentionally a **static surface paper**, not a dynamics paper. The claim is that today's surface can be represented, constrained, and reused more cleanly; modeling tomorrow's coefficient evolution is a natural next step, but it is outside the scope of the current manuscript.

1.4 Connection to the Spectral Fenton Framework

This paper extends the Spectral Fenton architecture (Nagy, 2026a,b) from 1D densities to 2D surfaces. The same Fourier cosine basis and the same coefficient-level reasoning — previously used to represent portfolio loss distributions — here parametrize the implied volatility surface. The arbitrage-free conditions become algebraic constraints on the cosine coefficients, verified by the same Lean 4 proof infrastructure.

That reuse is not only mathematically tidy; it is operationally valuable. A representation that is already compatible with COS pricing and coefficient-level risk objects is easier to turn into a desk tool than a standalone surface ansatz that stops at calibration.

1.5 Notation

Symbol	Meaning
$k = \log(K/F)$	Log-moneyness (forward)
$w(k, T)$	Total implied variance: $\sigma^2(K, T) \cdot T$
$C(K, T)$	European call price
$q(K)$	Risk-neutral density
r	Number of Fourier cosine modes
m	Number of maturity nodes
$u_j(T)$	j-th spectral coefficient at maturity T
ω_j	j-th frequency
$c(T)$	Constant (DC) spectral coefficient at maturity T

---

2. Preliminaries

2.1 Total Implied Variance

Following Gatheral (2004), we work with total implied variance $w(k, T) = \sigma^2_{\text{imp}}(K, T) \cdot T$ rather than implied volatility directly. The total variance is more natural for expressing no-arbitrage conditions, as it must be non-decreasing in maturity (calendar) and its curvature relates to the risk-neutral density (butterfly).

2.2 Calendar No-Arbitrage

A total variance surface $w(k, T)$ is free of calendar arbitrage if for all k and all $T_1 < T_2$:

$$w(k, T_2) \geq w(k, T_1).$$

Violation would permit a negative-cost calendar spread: sell the shorter-dated option, buy the longer-dated one.

2.3 Butterfly No-Arbitrage

The classical result of Breeden and Litzenberger (1978) states that the risk-neutral density is the second derivative of the call price with respect to strike:

$$q(K) = e^{rT} \frac{\partial^2 C}{\partial K^2}(K, T).$$

Since q is a probability density, it must be non-negative. Equivalently, $C(K, T)$ must be **convex** in K. Any violation implies a butterfly spread with negative cost. In the discrete case, this means:

$$C(K-\delta) + C(K+\delta) \geq 2\,C(K) \quad \text{for all } K, \delta > 0.$$

---

3. The Spectral Volatility Surface

3.1 Definition

We represent total implied variance as a finite Fourier cosine series:

$$w(k, T) = c(T) + \sum_{j=1}^{r} u_j(T) \cdot \cos(\omega_j\, k)$$

where: - $c(T) \in \mathbb{R}_+$ is the constant (DC) coefficient at maturity T; - $u_j(T) \in \mathbb{R}$ are the spectral coefficients ($j = 1, \dots, r$); - ω_j are the frequencies (typically $\omega_j = j\pi/k_{\max}$ for a domain $[-k_{\max}, k_{\max}]$); - r is the **rank** of the surface (number of Fourier modes per slice).

The cosine basis $\{\cos(\omega_j k)\}$ is natural for implied volatility: the surface is even in log-moneyness (approximately symmetric around the forward), and cosine functions are even.

3.2 Low Rank

In practice, implied volatility surfaces are smooth in log-moneyness. The Fourier cosine coefficients $u_j(T)$ decay rapidly with j, and $r = 6$–12 modes suffice for sub-basis-point accuracy on SPX surfaces. The total parameter count is $r \times m$ where m is the number of maturity nodes, yielding ~ 48–96 free parameters for a full surface — compared to $\sim 5m$ for SVI (without cross-maturity guarantees).

3.3 Maturity Interpolation

The spectral representation defines the surface at discrete maturity nodes $\{T_1, \dots, T_m\}$. Between nodes, the coefficient functions $c(T)$ and $u_j(T)$ can be interpolated — for instance via linear or monotone cubic interpolation — while preserving the Fejér constraints. Since the Fejér bound $c(T) \geq \sum_j |u_j(T)|$ is convex in the coefficients, linear interpolation between two feasible coefficient vectors remains feasible. This ensures that calendar no-arb holds not only at the maturity nodes but also at interpolated maturities, provided the interpolation scheme preserves convex combinations.

Figure 1. *3D surface plot: spectral total variance $w(k, T)$ versus log-moneyness and maturity, overlaid with market mid-quotes from 8 SPX maturities. The surface is smooth in k (cosine basis) and piecewise across maturities. Generated by examples/volsurface_figures.py.*

---

4. Arbitrage-Free Conditions

4.1 The Fejér Non-Negativity Condition

The algebraic core of our construction is a classical result from Fourier analysis:

Theorem 1 (Non-negative cosine polynomial). *Let* $c \in \mathbb{R}$, $\{a_j\}_{j=1}^r \subset \mathbb{R}$, *and* $\{\omega_j\}_{j=1}^r \subset \mathbb{R}$. *If*

$$c \geq \sum_{j=1}^{r} |a_j|,$$

then $c + \sum_{j=1}^{r} a_j \cos(\omega_j x) \geq 0$ *for all* $x \in \mathbb{R}$.

Proof. Since $|\cos(\omega_j x)| \leq 1$, we have $|a_j \cos(\omega_j x)| \leq |a_j|$ for each j. By the triangle inequality applied termwise:

$$\sum_{j=1}^{r} a_j \cos(\omega_j x) \geq -\sum_{j=1}^{r} |a_j| \geq -c.$$

Adding c to both sides: $c + \sum_j a_j \cos(\omega_j x) \geq 0$. □

Lean 4 verification: LeanProofs/VolSurface/NonNegCosine.lean — cosine_poly_nonneg_of_fejer. Compiles with 0 sorry.

This simple result — the constant term dominating the absolute values of the oscillatory coefficients — is the Fejér-type condition that powers all three no-arbitrage results.

4.2 Total Variance Non-Negativity

Theorem 2 (Spectral total variance is well-defined). *If the coefficients* $\{u_j(T)\}$ *satisfy the Fejér bound* $c(T) \geq \sum_{j=1}^{r} |u_j(T)|$ *for each maturity* T, *then* $w(k,T) \geq 0$ *for all* k, T.

Proof. Immediate from Theorem 1 applied pointwise for each T. □

Lean 4: LeanProofs/VolSurface/TotalVariance.lean — spectral_total_variance_nonneg.

4.3 Calendar No-Arbitrage

Theorem 3 (Monotone sum preservation). *If* $f_j \leq g_j$ *for all* $j = 1, \dots, r$, *then* $\sum_j f_j \leq \sum_j g_j$.

Lean 4: LeanProofs/VolSurface/MonotoneSum.lean — sum_mono_of_componentwise_mono.

Theorem 4 (Calendar no-arbitrage). *Let* $T_1 < T_2$ *and define the coefficient increments* $\Delta c = c(T_2) - c(T_1)$ *and* $\Delta u_j = u_j(T_2) - u_j(T_1)$. *If the increments satisfy the Fejér bound*

$$\Delta c \geq \sum_{j=1}^{r} |\Delta u_j|,$$

then $w(k, T_2) \geq w(k, T_1)$ *for all* k.

Proof. The difference $w(k, T_2) - w(k, T_1) = \Delta c + \sum_j \Delta u_j \cos(\omega_j k)$ is itself a cosine polynomial. The Fejér bound on the increments guarantees non-negativity (Theorem 1). □

Lean 4: LeanProofs/VolSurface/CalendarNoArb.lean — calendar_noarb_from_fejer_increment.

Remark. The calendar condition is a linear constraint on the coefficient functions: $\Delta c \geq \sum |\Delta u_j|$ is a polyhedral (specifically, ℓ^1-ball) constraint in coefficient space. This makes calibration subject to calendar no-arb a **convex constrained fit** — for example via projected least squares, linear-program surrogates, or SOCP formulations — rather than an unconstrained nonlinear search.

4.4 Butterfly No-Arbitrage

Theorem 5 (Cosine series second derivative). *The second derivative of the cosine series* $w(k) = c + \sum_j a_j \cos(\omega_j k)$ *is*

$$\frac{\partial^2 w}{\partial k^2}(k) = -\sum_{j=1}^{r} a_j \omega_j^2 \cos(\omega_j k).$$

Proof. By the chain rule: $\frac{d}{dk} \cos(\omega_j k) = -\omega_j \sin(\omega_j k)$, and $\frac{d^2}{dk^2} \cos(\omega_j k) = -\omega_j^2 \cos(\omega_j k)$. Linearity of differentiation extends to the finite sum. □

Lean 4: LeanProofs/VolSurface/CosineSecondDeriv.lean — cosine_series_first_deriv, cosine_series_second_deriv, and cosine_series_second_deriv_clean.

Theorem 6 (Breeden-Litzenberger discrete). *If* $C(K)$ *is convex in* K *(i.e.,* $C(K-\delta) + C(K+\delta) \geq 2C(K)$ *for all* $K, \delta > 0$*), then the risk-neutral density is non-negative.*

Proof. From convexity at $t = 1/2$: $C(\frac{1}{2}(K-\delta) + \frac{1}{2}(K+\delta)) \leq \frac{1}{2}C(K-\delta) + \frac{1}{2}C(K+\delta)$. Simplifying: $C(K) \leq \frac{1}{2}[C(K-\delta) + C(K+\delta)]$, hence $C(K-\delta) + C(K+\delta) - 2C(K) \geq 0$. □

Lean 4: LeanProofs/VolSurface/BreedenLitzenberger.lean — breeden_litzenberger_density_nonneg.

Theorem 7 (Non-negative curvature polynomial from spectral coefficients). *If the curvature-scaled coefficients satisfy the Fejér bound*

$$c_{\text{curv}} \geq \sum_{j=1}^{r} |a_j \omega_j^2|,$$

then the curvature polynomial $c_{curv} + \sum_j (a_j \omega_j^2) \cos(\omega_j k) \geq 0$ *for all* k.

Proof. The curvature-scaled coefficients $\{a_j \omega_j^2\}$ form a cosine polynomial. The Fejér bound guarantees non-negativity (Theorem 1). □

Remark. Non-negativity of the curvature polynomial is a *sufficient but not necessary* condition for the full Gatheral density $g(k) \geq 0$ (butterfly no-arb). The Gatheral density involves w, w', and w'' jointly in a nonlinear expression (see §2.3 and §13.2). Non-negative curvature ensures the $w''/2$ term in $g(k)$ is well-behaved, but the full condition requires the additional curvature-dominance analysis of Theorems L12–L14 (§12). The calendar no-arb guarantee (Theorem 4) is exact; the butterfly guarantee via Theorem 7 alone is partial.

Lean 4: LeanProofs/VolSurface/ButterflyNoArb.lean — butterfly_noarb_curvature_nonneg and butterfly_noarb_from_curvature.

Figure 2. *Arbitrage violation heatmap: Spectral surface (left) vs. raw SVI (right) on a stressed SPX surface (VIX ≈ 30). Strike (x-axis) vs. maturity (y-axis), colored by Gatheral $g(k)$ violation magnitude. The spectral surface shows zero or near-zero violations across the grid; raw SVI shows 100+ violations concentrated at short maturities and far out-of-the-money strikes. Generated by examples/volsurface_figures.py.*

---

5. Combining the Conditions

Theorems 2, 4, and 7 establish three independent Fejér conditions on the spectral coefficients — for non-negative variance, calendar no-arb, and non-negative curvature, respectively. Their conjunction, assembled in LeanProofs/VolSurface/ArbitrageFree.lean (spectral_surface_arbitrage_free), gives the base arbitrage-free result. We defer the full statement to the capstone theorem (Theorem 10, §8), which additionally incorporates the parameter count (§6) and COS pricing (§7).

---

6. Parameter Count and Feasible Set

Theorem 8 (Finite parameter count). *For r basis functions and m maturity nodes, the spectral surface has exactly $r \cdot m$ free parameters. The feasible set (satisfying all three Fejér conditions) is a convex polytope in $\mathbb{R}^{r \cdot m}$.*

Proof. The parameter space is $\{u_j(T_i)\}$ for $j \in \{1, \dots, r\}$, $i \in \{1, \dots, m\}$, giving $r \cdot m$ parameters. Each Fejér condition is a linear inequality of the form $c \geq \sum |u_j|$, which defines a polyhedron (ℓ^1-ball constraint) in coefficient space. The intersection of polyhedra is a convex polytope. ☐

Lean 4: LeanProofs/VolSurface/ParameterCount.lean — spectral_surface_param_count.

Remark. The convexity of the feasible set is crucial for calibration: fitting the spectral surface to market data subject to no-arb constraints is a **convex optimization problem**, solvable globally in polynomial time via linear programming or second-order cone programming (SOCP).

---

7. COS Pricing from the Spectral Surface

Theorem 9 (COS pricing from spectral coefficients). *Given density coefficients $\{A_k\}$ determined by the spectral surface and payoff coefficients $\{V_k\}$ (analytic for European calls/puts), the option price is:*

$$C(K, T) = e^{-rT} \sum_{k=0}^{N-1} A_k \cdot V_k.$$

If the density coefficients decompose spectrally as $A_k = \sum_j w_j \cdot B_{jk}$, then the price admits the factored form:

$$C(K,T) = e^{-rT} \sum_{j=1}^{r} w_j \left(\sum_{k=0}^{N-1} B_{jk} \cdot V_k \right).$$

No numerical integration or simulation is required.

Proof. The first identity is the standard COS pricing formula (Fang and Oosterlee, 2008). The factored form follows by swapping the order of summation:

$$\sum_k \left(\sum_j w_j B_{jk} \right) V_k = \sum_j w_j \sum_k B_{jk} V_k.$$

This is valid for finite sums. □

Lean 4: LeanProofs/VolSurface/COSPricing.lean — cos_pricing_from_spectral_coeffs.

Remark. The factored form has computational significance: the inner sums $\sum_k B_{jk}V_k$ can be precomputed once per payoff type, making subsequent repricing $O(r)$ — a constant-time operation in the number of strikes and maturities.

---

8. Main Theorem: The Spectral Volatility Surface

Theorem 10 (Spectral Volatility Surface — Capstone). *Let $w(k,T) = c(T) + \sum_j u_j(T)\cos(\omega_j k)$ be a spectral volatility surface with r modes per maturity and m maturity nodes. If the coefficient functions $\{c(T), u_j(T)\}$ satisfy the three Fejér conditions:*

1. *Fejér bound at each maturity: $c(T) \geq \sum_j |u_j(T)|$ (Theorem 2 — total variance non-negativity);*
2. *Fejér bound on increments: $\Delta c \geq \sum_j |\Delta u_j|$ for consecutive maturities (Theorem 4 — calendar no-arb);*
3. *Fejér bound on curvature coefficients: $c_{curv} \geq \sum_j |u_j \omega_j^2|$ (Theorem 7 — non-negative curvature polynomial);*

then the surface is simultaneously:

- *(a) A valid total variance surface: $w(k,T) \geq 0$ for all k, T;*
- *(b) Calendar-arbitrage-free by construction;*
- *(c) Non-negative in curvature (a sufficient condition for butterfly no-arb; see §4.4 Remark);*
- *(d) Parametrized by exactly $r \cdot m$ coefficients in a convex feasible set (Theorem 8);*
- *(e) Equipped with closed-form COS pricing: $C(K,T) = e^{-rT} \sum_j w_j (\sum_k B_{jk} V_k)$ (Theorem 9).*

Proof. Properties (a)–(c) follow from the conjunction of Theorems 2, 4, and 7 (assembled in ArbitrageFree.lean). Property (d) follows from Theorem 8 (parameter count and convex polytope structure). Property (e) follows from Theorem 9 (COS pricing factorization). The Lean capstone (MainTheorem.lean) verifies the conjunction. □

Lean 4: LeanProofs/VolSurface/MainTheorem.lean — spectral_vol_surface_main. Compiles with 0 sorry.

Remark. Theorem 10 unifies the entire construction: the spectral basis provides the representation (§3), the Fejér conditions provide the guarantees (§4), the parameter count provides tractability (§6), and the COS factorization provides pricing (§7). A single set of $r \times m$ spectral coefficients simultaneously defines the surface, guarantees no-arb, and prices options — with no auxiliary computation.

---

9. Dupire Local Volatility from Spectral Coefficients

The spectral representation yields **closed-form Dupire local volatility**. Since $w(k) = c + \sum_j u_j \cos(\omega_j k)$, the derivatives are explicit:

$$w'(k) = -\sum_j u_j \omega_j \sin(\omega_j k), \qquad w''(k) = -\sum_j u_j \omega_j^2 \cos(\omega_j k)$$

The maturity derivative $\partial w/\partial T$ is computed from the coefficient functions $c(T), u_j(T)$ by finite differences or interpolation. Dupire's formula then gives:

$$\sigma_{\text{loc}}^2(k,T) = \frac{\partial w/\partial T}{g(k)} \quad \text{where} \quad g(k) = \left(1 - \frac{kw'}{2w}\right)^2 - \frac{(w')^2}{4}\left(\frac{1}{w} + \frac{1}{4}\right) + \frac{w''}{2}$$

No numerical differentiation is needed: w, w', w'' are all analytic in the coefficients. The Gatheral density $g(k)$ in the denominator is the same quantity appearing in the butterfly condition (Section 4.4). When $g(k) > 0$ (ensured by the curvature-dominance condition L14), the local vol is well-defined and positive.

Practical significance: A derivatives desk can extract the full local vol grid — needed for exotic option pricing via PDE methods — directly from the $r \times m$ spectral coefficients, without any numerical derivative computation.

---

10. Spectral Greeks

All Black-Scholes Greeks are closed-form functions of the spectral coefficients. Given $\sigma(k) = \sqrt{w(k)/T}$ and the standard BS formulas:

- **Delta, Gamma, Vega, Theta** follow from BS with the spectral implied vol.
- **Vanna** ($\partial\text{Delta}/\partial\sigma$) and **Volga** ($\partial\text{Vega}/\partial\sigma$) incorporate $\partial\sigma/\partial k = w'/(2\sigma T)$, which is explicit from the coefficients.
- **Coefficient sensitivity:** $\partial C/\partial u_j = \text{Vega} \cdot \partial\sigma/\partial u_j = \text{Vega} \cdot \cos(\omega_j k)/(2\sigma T)$. This tells you how much each spectral mode contributes to the option price — a spectral decomposition of the smile risk.

Lean reuse: GreeksLinearity.lean (delta/gamma linearity) and VegaFormula.lean (Vega > 0) from the SpectralFenton directory apply directly. The spectral coefficient sensitivity is a new result specific to the vol surface representation.

---

11. Feasible Set Geometry

The Fejér constraints $c \geq \sum |u_j|$ define an ℓ^1-ball (diamond/cross-polytope) in coefficient space. For $r = 2$ modes, the feasible set $\{(u_1, u_2) : |u_1| + |u_2| \leq c\}$ is a diamond centered at the origin with vertices at $(\pm c, 0)$ and $(0, \pm c)$.

The calendar constraint $\Delta c \geq \sum |\Delta u_j|$ restricts the *second* maturity's coefficients to lie within a shifted diamond centered at the first maturity's coefficients. The intersection of the Fejér and calendar diamonds is the joint feasible set — still a convex polytope, hence tractable for optimization.

Key property: convexity of the feasible set guarantees that any convex combination of feasible coefficient vectors is also feasible. This means blending or interpolating between calibrated surfaces preserves all no-arb conditions — a property that neural network approaches cannot guarantee.

Figure 3. *Fejér feasible set for $r = 2$ modes. The constraint $|u_1| + |u_2| \leq c$ defines a diamond (cross-polytope) in the (u_1, u_2) plane, centered at the origin with vertices at $(\pm c, 0)$ and $(0, \pm c)$. The calendar constraint further restricts the second maturity to a shifted diamond centered at the first maturity's coefficients. The joint feasible set (intersection, shaded) is a convex polytope. Generated by examples/volsurface_figures.py.*

---

12. Formal Verification

All 63 theorems are formalized and machine-verified in Lean 4 (version 4.28.0) using Mathlib. The proof files reside in LeanProofs/VolSurface/ with the following structure:

Level	Theorem	File	Dependencies
L01	Non-negative cosine polynomial	NonNegCosine.lean	—
L02	Monotone sum preservation	MonotoneSum.lean	—
L03	Cosine second derivative	CosineSecondDeriv.lean	—
L04	Breeden-Litzenberger density	BreedenLitzenberger.lean	—
L05	Spectral total variance ≥ 0	TotalVariance.lean	L01
L06	Calendar no-arbitrage	CalendarNoArb.lean	L01, L02
L07	Butterfly no-arb (Fejér curvature)	ButterflyNoArb.lean	L01, L03, L04
L08	Spectral surface arb-free	ArbitrageFree.lean	L05, L06, L07
L09	Finite parameter count	ParameterCount.lean	—
L10	COS pricing from coefficients	COSPricing.lean	—
L11	Main theorem	MainTheorem.lean	L08, L09, L10

Level	Theorem	File	Dependencies
L12	**Gradient bound (sine)**	GradientBound.lean	—
L13	**Explicit margin** ($w \geq \delta$)	ExplicitMargin.lean	L01
L14	**Gatheral sufficient condition**	GatheralBound.lean	—
L15	**Chebyshev-cosine duality**	ChebyshevDuality.lean	L01

L12–L14: Narrowing the butterfly gap. Theorems L12–L14 address the gap between the Fejér curvature condition (L07) and the full Gatheral density $g(k) \geq 0$. The chain is:

- **L12:** $\left|\sum a_j \sin(\omega_j k)\right| \leq \sum |a_j|$ — bounds the derivative of a cosine series, giving $|w'(k)| \leq G$ where $G = \sum |u_j \omega_j|$.
- **L13:** If $c \geq \sum |u_j| + \delta$ with $\delta > 0$, then $w(k) \geq \delta$ — the total variance is bounded away from zero, preventing the Gatheral denominator from degenerating.
- **L14:** If $4\delta \cdot w'' \geq G^2 \cdot (4 + w_{\max})$ where δ is the Fejér margin, G^2 is the squared gradient bound, and $w_{\max}$ is the maximum total variance, then $g(k) \geq 0$ — the curvature dominates the gradient-squared term, guaranteeing the full Gatheral density condition.

Together, L12–L14 provide an algebraic sufficient condition for the **full** Gatheral butterfly no-arb, expressed in terms of spectral coefficient bounds. This condition is satisfied on moderate-skew surfaces (SSVI-class) but may fail on high-skew configurations ($\rho < -0.5$). See §13.2.

Verification: 15 files, 63 theorems, 0 sorry, 0 errors. All proofs use only Lean's core type theory and Mathlib's real analysis formalization.

---

13. Discussion

13.1 Comparison with Existing Methods

Property	SVI/SSVI	Heston	Neural	MR-SABR	Cheby. Tensor	SANOS	**Spectral (ours)**
Params/slice	~ 5	implicit	$\sim 10^3$	~ 5	many	many	r (6–12)
Low-rank	Yes	—	No	Yes	No	No	**Yes**
Calendar no-arb	No	Yes	No	Yes	Yes	Yes	**Yes**
Butterfly no-arb	No	Yes	No	Yes	Yes	Yes	**Partial***
Model-free	Yes	No	Yes	No	Yes	Yes	**Yes**
COS pricing	No	No	No	No	No	No	**Yes**

Property	SVI/SSVI	Heston	Neural	MR-SABR	Cheby. Tensor	SANOS	**Spectral (ours)**
Shared risk infra	No	No	No	No	No	No	**Yes**
Machine-verified	No	No	No	No	No	No	**Yes**

The mean-reverting SABR (MR-SABR) is the closest competitor on parameter count and no-arb guarantees, but it is model-specific (tied to SABR dynamics) and does not share infrastructure with pricing or risk measurement. The Chebyshev tensor approach is closest in mathematical spirit (polynomial basis + constraints) but uses a different basis, does not achieve low rank, and lacks formal verification. SANOS guarantees no-arb via linear programming but is non-parametric (many parameters).

*Partial: Fejér guarantees non-negative curvature polynomial, which is a sufficient but strictly weaker condition than the full Gatheral $g(k) \geq 0$. The strengthened condition L12–L14 partially closes this gap — see §13.2 for the precise scope.

Our unique combination: **low-rank + model-free + calendar-arb-free + COS pricing + unified risk infrastructure + machine-verified proofs**. No other method in the literature achieves all six simultaneously.

13.2 Honest Limitations

We identify three limitations that a reader should understand before applying this framework.

Limitation 1: The Fejér bound is sufficient but far from necessary. The condition $c \geq \sum |a_j|$ guarantees $c + \sum a_j \cos(\omega_j x) \geq 0$, but many non-negative cosine polynomials violate this bound. The necessary and sufficient condition requires a sum-of-squares (SOS) decomposition, expressible as a semidefinite program (SDP). In our Monte Carlo experiments (50 random SSVI surfaces, $r = 6$), 100% of unconstrained least-squares fits were already Fejér-feasible, and the tightness ratio (constrained RMSE / unconstrained RMSE) was 1.0 — suggesting the Fejér bound is not binding on smooth, SSVI-class surfaces. However, for surfaces with high curvature or large skew (e.g., Heston-generated with $\rho = -0.7$), the bound becomes restrictive.

Limitation 2: Fejér does not guarantee the full Gatheral butterfly condition. The Gatheral (2004) density condition for butterfly no-arbitrage is:

$$g(k) = \left(1 - \frac{kw'}{2w}\right)^2 - \frac{(w')^2}{4}\left(\frac{1}{w} + \frac{1}{4}\right) + \frac{w''}{2} \geq 0$$

This involves w, w', and w'' jointly — a nonlinear condition. Our Lean proofs verify that the Fejér bound ensures the curvature cosine polynomial is non-negative, which is a **weaker** condition than $g(k) \geq 0$. On Heston-calibrated surfaces ($\rho = -0.7$), we observe 46 Gatheral violations with Fejér constraints — reduced from 134 (raw SVI) and 312 (SSVI), indicating the Fejér structure provides substantial but incomplete butterfly protection. The calendar no-arb guarantee is exact; the butterfly guarantee is approximate. Closing this gap — proving coefficient conditions that imply the full Gatheral density condition — is the primary open problem.

Limitation 3: Real-data validation is still a first-pass snapshot, not a full panel benchmark. Our strongest controlled evidence still comes from SSVI-generated (Gatheral-Jacquier 2014) and Heston-calibrated surfaces. We now also include a staged real-SPX common-grid sanity check, but that is narrower than a production-grade panel fit: it uses 8 maturities, one shared log-moneyness interval, and interpolated grid metrics rather than a full raw-quote objective across the whole chain. Validation on richer real SPX/EuroStoxx panels remains the main empirical next step.

13.3 Benchmark Summary

Data source	Method	RMSE	Calendar viol.	Butterfly (Gatheral)	Params
SSVI + noise	**Spectral** (r=6)	0.011	**0**	**0**	56
SSVI + noise	SSVI	0.001	0	0	11
SSVI + noise	Raw SVI	0.013	47	23	40
Heston (ρ=-0.7)	**Spectral** (r=6)	0.386	**0**	46	56
Heston (ρ=-0.7)	SSVI	1.446	0	312	11
Heston (ρ=-0.7)	Raw SVI	0.149	151	134	40

On SSVI-class data, the spectral surface performs well: calendar-arb-free by construction, zero Gatheral violations, competitive RMSE.

Figure 4. *Grouped bar chart: RMSE (top) and violation counts (bottom) across methods and data sources. Left panel: SSVI + noise. Right panel: Heston ($\rho = -0.7$). The spectral surface achieves zero calendar violations in both cases; the Cheby+SDP variant achieves the lowest RMSE. Generated by examples/volsurface_figures.py.*

SPX-Realistic Benchmark (3 Market Regimes)

We test on Heston-calibrated surfaces matching published SPX parameters (Gatheral, 2006) across Normal (VIX\$ 16),$Stressed(VIX\ 30),andCalm(VIX$ \$12) regimes. Maturities: 1W–2Y. Strikes: 50%–150% moneyness. Butterfly checked via real Gatheral density $g(k)$.

Regime	Method	RMSE	Cal.	But. (Gatheral)	Params
Normal	**Cheby+SDP (r=8)**	**0.0036**	6	9	54
Normal	Cosine+Fejér (r=12)	0.0832	**0**	2	78
Normal	Raw SVI	0.0289	146	84	30
Stressed	**Cheby+SDP (r=8)**	**0.0029**	3	**0**	54
Stressed	Cosine+Fejér (r=12)	0.1256	0	8	78
Stressed	Raw SVI	0.0253	119	95	30

Regime	Method	RMSE	Cal.	But. (Gatheral)	Params
Calm	**Cheby+SDP (r=8)**	**0.0026**	5	**0**	54
Calm	Cosine+Fejér (r=12)	0.0357	0	4	78
Calm	Raw SVI	0.0131	119	114	30

Key findings: (1) In the released fixed-seed benchmark, Cheby+SDP achieves RMSE in the narrow range 0.0026–0.0036 across all three SPX-like regimes. (2) On the **stressed** and **calm** regimes, Cheby+SDP achieves **zero Gatheral butterfly violations**; on the normal regime it still leaves a small residual count. (3) Cosine+Fejér with $r = 12$ modes achieves 0 calendar violations in every regime and keeps butterfly violations small (2–8), which makes it the cleanest strict calendar-control layer. (4) Raw SVI produces 119–146 calendar violations per regime together with 84–114 butterfly violations, confirming why cross-maturity guarantees matter.

The remaining Cheby+SDP calendar violations (3, 5, and 6 in the three reported regimes) appear to be implementation-level solver-tolerance effects rather than failures of the underlying constrained representation. In the current evidence surface, the cosine Fejér layer is the one that enforces exact zero calendar violations most cleanly.

Observed Calibration Time

The released benchmark currently reports full-surface calibration time directly, and these are the numbers that should anchor the paper:

Method	Normal	Stressed	Calm
Cheby+SDP (r=8)	230.7 ms	54.4 ms	91.9 ms
Cosine+Fejér (r=8)	1.1 ms	1.0 ms	1.0 ms
Cosine+Fejér (r=12)	1.3 ms	1.3 ms	1.2 ms
Raw SVI	131.6 ms	55.6 ms	47.5 ms

These timings come from the fixed-seed run of examples/spx_benchmark.py reported in this draft and should be read as implementation-specific wall-clock observations rather than hardware-independent asymptotic claims. The desk-relevant conclusion is still strong: the direct cosine layer calibrates in about a millisecond, while the Chebyshev refinement remains sub-second even with solver overhead.

The representation also supports analytic repricing, Greeks, and local-vol extraction from the same coefficient object, but separate runtime tables for those downstream operations are not yet part of the released benchmark harness. We therefore treat those downstream speed claims as structural consequences of the representation rather than as standalone timing evidence in this draft.

Reproducibility. The empirical results claimed in this section are supported by the released scripts examples/spx_benchmark.py, src/spectral_fenton/vol_surface.py, and src/spectral_fenton/vol_surface_cheby.py. The benchmark runner now supports:

- the fixed-seed synthetic SPX-realistic regime suite, and
- a staged real-data snapshot benchmark from data/equities/spx_real_options.csv, and
- a separate zero-cost historical SPX archive lab built from normalized public optiondata.org data

via ——dataset synthetic, ——dataset real, and ——dataset optiondata—free—aggregate—validate, respectively.

Staged Real SPX Snapshot (Common-Grid Sanity Check)

As a first real-data sanity check, we take the staged SPX options snapshot (11,275 call quotes across 50 expiries), select 8 representative maturities from roughly 1M to 1.8Y, intersect them to a common log-moneyness interval, and interpolate each maturity slice onto an 80-point shared grid. This is intentionally narrower than a full raw-quote panel fit: it is a clean cross-maturity comparison surface rather than the final production benchmark.

Method	RMSE	Calendar viol.	Butterfly (Gatheral)	Params
Cheby+SDP (r=8)	0.000004	0	1	72
Cosine+Fejér (r=8)	**0.000001**	**0**	**0**	72
Cosine+Fejér (r=12)	**0.000001**	**0**	**0**	104
SSVI	0.000063	0	0	11
Raw SVI	0.047746	80	0	40

This real-snapshot benchmark is revealing in a different way from the stressed synthetic regime tests. On the shared grid, the staged SPX snapshot is smooth enough that the spectral fits and SSVI all remain structurally clean, while raw SVI still breaks badly at the calendar level. The strongest practical discriminator in this first real-data harness is therefore not butterfly cleanup but **cross-maturity coherence plus fit quality under one reusable surface object**.

At the same time, this should not be oversold. The common-grid snapshot is a **reviewer-safe first real-data benchmark**, not the final desk benchmark. The next empirical step is still the fuller panel fit: use more maturities, preserve more of the wings, and report raw-quote rather than common-grid metrics alongside the same arbitrage diagnostics.

Figure 5. *Real SPX snapshot smile slices on a common grid: market mid-surface (dots) versus Cheby+SDP, Cosine+Fejér, SSVI, and raw SVI at short, medium, and long maturities. Generated by examples/spx_benchmark.py ——dataset real.*

Zero-Cost Historical SPX Robustness Lab

In addition to the staged modern snapshot, the released harness now includes a separate **zero-cost historical SPX robustness lab** built from the public Jan-Jun 2013 optiondata.org archive. These data are normalized into a grouped local file, data/equities/spx_free_optiondata_2013H1.csv, covering 124 trading dates and roughly 370k option rows. Methodologically, this archive serves a different purpose from the staged modern snapshot: it is a cheap robustness and engineering laboratory, not the reviewer-facing evidence surface for current-market claims.

That distinction matters. The staged modern snapshot is the place where this draft makes its current-market real-data claim: on a clean shared grid, the reusable spectral surface remains structurally clean and fits very accurately. The free 2013 archive is instead where we ask a harsher engineering question: if we group the archive by quote_date and validate the same family across many dates, how stable is the surface architecture under older regimes and noisier panels?

On the full grouped 2013H1 run, the unguarded practical spectral preset (scale=0.35, rank 6, wing_linear) is already strong: it reaches 120/124 stable dates with mean RMSE 0.009762. A narrow per-slice butterfly repair layer improves this further without changing the global panel preset: the guarded practical preset reaches **124/124 stable dates, zero total butterfly violations, and mean RMSE 0.009761**. By contrast, SSVI still has the best average RMSE on this older archive (0.002536) but shows one severe outlier day with a large butterfly-violation cluster.

The right conclusion is therefore not that the free archive replaces the staged modern benchmark. It does not. The right conclusion is that the spectral surface now has a credible **zero-cost historical robustness story** as well: on an older multi-date SPX archive, a low-rank wing-mapped fit plus a small local repair layer can stay arbitrage-clean across the full grouped run while preserving essentially the same fit quality. That is an engineering result first, and it should be read as such.

The grouped archive also gives a first view of **coefficient stability**, not just fit stability. On the full 124-day run, the two guarded spectral variants are quite similar in coefficient-space behavior: neither achieves its robustness by introducing visibly wilder coefficient motion than the other. The main coefficient instability is localized rather than global, and it concentrates primarily in the shortest tenor bucket. Figure 6 makes that localization visible directly across tenor buckets: the front bucket is the dominant hotspot, the long end is the next-most mobile region, and the middle of the curve is materially calmer.

Guarded method	Mean active rank	Mean adjacent coefficient drift	Worst-slice 90th-percentile jump	Most unstable bucket
Free practical + guard	5.84	0.000607	0.005042	Front bucket (about 30D)
Modern local default + guard	5.81	0.000605	0.004742	Front bucket (about 30D)

This matters for interpretation. The old-archive weight regime switches are real, but they do not imply that the underlying guarded spectral representation is unstable everywhere. The stronger reading is more local: the front tenor remains the most sensitive coefficient bucket, while the middle of the curve is materially calmer.

Figure 6. *Guarded spectral coefficient stability by tenor on the full grouped free SPX archive. Left to right: median adjacent coefficient drift, 90th-percentile adjacent coefficient drift, and mean active rank for the two guarded spectral variants. The shortest bucket around 30D is the dominant local hotspot, while the middle of the curve is materially calmer. Generated by examples/spx_benchmark.py −−dataset optiondata−free−stability−plot.*

13.4 What the Lean Proofs Actually Prove

To be precise about what is and is not machine-verified:

Claim	Verified?	Lean file
Fejér ⇒ non-negative cosine polynomial	**Yes**	NonNegCosine.lean
Calendar no-arb from coefficient monotonicity	**Yes**	CalendarNoArb.lean

Claim	Verified?	Lean file
Curvature coefficients ⇒ non-negative curvature	**Yes**	ButterflyNoArb.lean
$\lvert w'(k)\rvert \le G$ (gradient bound)	**Yes**	GradientBound.lean
$w(k) \ge \delta > 0$ (explicit margin)	**Yes**	ExplicitMargin.lean
Curvature dominance ⇒ Gatheral $g(k) \ge 0$	**Yes**	GatheralBound.lean
COS pricing as finite sum	**Yes**	COSPricing.lean

The butterfly gap is now **partially closed**: L14 proves that the Gatheral density $g(k) \ge 0$ when the curvature is large relative to the squared gradient and the total variance is bounded away from zero. The remaining gap is that this "curvature dominance" condition may fail for high-skew surfaces where $|w'|$ is large relative to w''. The calendar no-arb is fully verified with no gap.

13.5 Chebyshev Variant: SDP-Enforced Constraints

The Chebyshev polynomials $T_n(x)$ are dual to the cosine basis via $T_n(\cos\theta) = \cos(n\theta)$. This means the spectral volatility surface can equivalently be represented as:

$$w(k,T) = \sum_{j=0}^{r} c_j(T)\, T_j(k/k_{\max})$$

with the same spectral coefficients. The key advantage: for polynomial non-negativity (and by extension, the Gatheral density), **semidefinite programming (SDP)** can enforce constraints exactly rather than relying on the loose Fejér bound.

We implement a Chebyshev variant (ChebyshevVolSurface) that solves a QP per maturity slice:

$$\min_{c} \|w_{\text{fit}} - w_{\text{market}}\|^2 \quad \text{s.t.} \quad w(k_i) \ge 0, \quad w(k_i,T) \ge w(k_i,T_{\text{prev}})$$

using CLARABEL (a modern interior-point SDP solver). Results:

Data	Method	RMSE	Cal.	But. (Gatheral)	Params
SSVI	**Cheby+SDP**	**0.001**	**0**	**0**	56
SSVI	Cosine+Fejér	0.011	0	0	56
Heston	**Cheby+SDP**	**0.002**	**1**	26	56
Heston	Cosine+Fejér	0.387	0	16	56
Heston	Raw SVI	0.150	148	130	40

On SSVI data, Chebyshev+SDP achieves 10x better RMSE than Cosine+Fejér with identical zero violations. On Heston data ($\rho = -0.7$), Chebyshev achieves **200x better RMSE** (0.002 vs 0.387) with near-zero calendar violations. The butterfly violations (26 vs 16) reflect the different trade-off:

the cosine approach squashes the fit to avoid violations (at the cost of accuracy), while Chebyshev preserves the surface shape.

The Chebyshev-cosine duality is verified in Lean 4 (ChebyshevDuality.lean, L15): the Fejér bound, calendar condition, and COS pricing formula all carry over. 15 theorems total.

Recommendation: For maximum accuracy, use Chebyshev+SDP. For guaranteed calendar no-arb with minimal implementation complexity, use Cosine+Fejér. Both share the same COS pricing infrastructure.

13.6 Open Questions

1. **Closing the butterfly gap.** The key open problem. Can one derive algebraic coefficient conditions (Fejér-type or SOS) that guarantee $g(k) \geq 0$ directly? If so, can these be machine-verified?
2. **SOS relaxation.** The SOS non-negativity condition is necessary and sufficient for polynomial non-negativity (Lasserre, 2001; Blekherman, Parrilo, and Thomas, 2013), but requires SDP solvers. Can real-time calibration (< 100 ms) be achieved with modern SDP solvers (MOSEK, SCS)?
3. **Minimum rank for realistic surfaces.** What is the smallest r that achieves sub-basis-point accuracy on real market data?
4. **Local volatility extraction.** Dupire's formula has a closed form in spectral coefficients. How does the Fejér constraint interact with the denominator's positivity?
5. **Stochastic vol connection.** Does Heston's characteristic function decompose naturally in the cosine basis?

---

14. Conclusion

We have introduced the spectral volatility surface: a low-rank Fourier cosine representation of total implied variance with algebraic coefficient constraints that partially guarantee arbitrage-free construction. The framework achieves:

- **Low rank:** $O(r \cdot m)$ parameters ($\sim$ 48–96 for a full surface).
- **Calendar no-arbitrage by construction:** Fejér bound on coefficient increments, machine-verified in Lean 4.
- **Butterfly no-arb (partial):** Fejér bound ensures non-negative curvature polynomial, a sufficient but not necessary condition for the full Gatheral density condition. On smooth surfaces (SSVI-class), this suffices empirically; on high-skew surfaces, a gap remains.
- **Analytic pricing:** European option prices via COS-method inner products.
- **Machine-verified proofs:** 63 Lean 4 theorems, 0 sorry. To our knowledge, the first formal verification of volatility surface no-arb conditions, including a sufficient condition for the full Gatheral density and the Chebyshev-cosine duality.
- **Basis-agnostic:** Both Fourier cosine (for simplicity) and Chebyshev polynomial (for SDP-enforced accuracy) bases are supported, connected via $T_n(\cos\theta) = \cos(n\theta)$.

The calendar result is exact. The butterfly result is two-layered: the base Fejér condition (L07) guarantees non-negative curvature; the strengthened condition (L12–L14) guarantees the full Gatheral density under a curvature-dominance assumption. On smooth surfaces (SSVI-class), both conditions are satisfied. For high-skew surfaces, the remaining gap is between curvature dominance and the full nonlinear Gatheral condition — the primary target for future work.

Acknowledgements

The Lean 4 proofs were developed with assistance from the Mathlib community. AI tools (Claude, Cursor) were used for drafting, code generation, and proof exploration. All mathematical results were independently verified via the Lean proof assistant; no theoretical claim relies solely on AI-generated output. The author thanks the anonymous reviewers for helpful comments.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Ackerer, D., Tagasovska, N., and Vatter, T (2020). Deep smoothing of the implied volatility surface. *Proceedings of NeurIPS.*.
- Andersen, L. and Brotherton-Ratcliffe, R (1997). The equity option volatility smile: An implicit finite-difference approach. *Journal of Computational Finance*, 1(2), 5-37.
- Blekherman, G., Parrilo, P.A., and Thomas, R.R (2013). Semidefinite Optimization and Convex Algebraic Geometry. *Semidefinite Optimization and Convex Algebraic Geometry.*
- Breeden, D. T. and Litzenberger, R. H (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4). DOI: 10.1086/296025
- Cohen, S.N., Reisinger, C., and Wang, S (2023). Arbitrage-free neural-SDE market models. *Mathematical Finance*, 33(2), 362-405.
- Fang, Fang and Oosterlee, Cornelis W. (2008). A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions. *SIAM Journal on Scientific Computing*, 31(2), 826-848. DOI: 10.1137/080718061
- De Marco, S. and Martini, C (2018). Calibration of the SSVI model with no arbitrage condition. *Risk.*
- Fejer, L (1915). Uber trigonometrische Polynome. *Journal fur die reine und angewandte Mathematik.*
- Gatheral, J (2004). A parsimonious arbitrage-free implied volatility parameterization with application to the valuation of volatility derivatives.
- Gatheral, J (2006). The Volatility Surface: A Practitioner's Guide. *The Volatility Surface: A Practitioner's Guide.*

- Gatheral, J. and Jacquier, A (2014). Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1). DOI: 10.2139/ssrn.2033323
- Gatheral, J., Jaisson, T., and Rosenbaum, M (2018). Volatility is rough. *Quantitative Finance*, 18(6), 933-949.
- Heston, Steven L. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies*, 6(2), 327-343. DOI: 10.1093/oso/9780199257195.003.0014
- Lasserre, J.B (2001). Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3), 796-817.
- Lee, R (2004). The moment formula for implied volatility at extreme strikes. *Mathematical Finance*, 14(3), 469-480.
- Nagy, T. (2026). The Fenton Distribution Solved. *Working paper.*
- Nagy, T. (2026). Exact Portfolio VaR Without Monte Carlo: The Eigen-COS Method. *Zenodo.* DOI: 10.5281/zenodo.18910516
- Gatheral, J. and Jacquier, A (2014). Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1). DOI: 10.2139/ssrn.2033323
- Vladimirov, E (2023). iCOS: Option-Implied COS Method.
- Glau, K., Kressner, D., and Statti, F (2025). “Low-rank tensor approximation of implied volatility surfaces.” arXiv:2512.01967.
- Ackerer, D. and Filipovic, D (2026). “SANOS: Semi-Analytic Non-parametric Option-implied Surfaces.” arXiv:2601.11209.
- Cui, Z., Kirkby, J.L., and Nguyen, D (2024). Randomized SVI. *Quantitative Finance.*
- Cuchiero, C., Khosrawi, M., and Teichmann, J (2025). Mean-reverting SABR volatility surfaces. *Mathematical Finance.*

Appendix A: Lean Proof Summaries

L01 (NonNegCosine.lean). The proof uses abs_le to bound $|\cos(\omega x)| \leq 1$, abs_mul to factor $|a\cos(\omega x)| = |a| \cdot |\cos(\omega x)|$, Finset.sum_le_sum to lift the termwise bound to the sum, and linarith for the final arithmetic. 12 lines of tactic proof.

L02 (MonotoneSum.lean). Direct application of Finset.sum_le_sum. One-line proof.

L03 (CosineSecondDeriv.lean). Chain rule for cosine: hasDerivAt_id → .const_mul → .congr_deriv (mul_one) → .cos for the first derivative; .sin.neg.mul_const for the second. HasDerivAt.sum with a function-form conversion via Finset.sum_apply lifts to finite sums.

L04 (BreedenLitzenberger.lean). Applies ConvexOn.2 with weights $a = b = 1/2$, simplifies $\frac{1}{2}(K - \delta) + \frac{1}{2}(K + \delta) = K$ via ring, and closes with linarith.

L06 (CalendarNoArb.lean). Applies L01 to the coefficient increments $\Delta u_j = u_2(j) - u_1(j)$, then uses ring to expand $(u_2 - u_1) \cdot \cos = u_2 \cdot \cos - u_1 \cdot \cos$ and Finset.sum_sub_distrib to split the sum.

L08 (ArbitrageFree.lean). Pure conjunction: L05, L05, L06, L07.

L11 (MainTheorem.lean). Pure conjunction of L08, L09, L10.

L12 (GradientBound.lean). The sine analogue of L01: $|a\sin(\omega x)| \leq |a|$ since $|\sin| \leq 1$. Uses abs_le.mpr with neg_one_le_sin and sin_le_one, then Finset.sum_le_sum for both upper and lower bounds, combined via abs_le.mpr.

L13 (ExplicitMargin.lean). From L01 applied with $c - \delta$ as the constant term: if $c - \delta \geq \sum |a_j|$, then $(c - \delta) + \sum a_j \cos(\omega_j k) \geq 0$, hence $c + \sum a_j \cos(\omega_j k) \geq \delta$. Two lines of tactic proof.

L14 (GatheralBound.lean). The key new result. Multiplied-out form avoids division: if $4\delta \cdot w'' \geq G^2(4 + w_{\max})$, then $4\delta \cdot w'' \geq (w')^2(4 + w_{\max})$ (since $(w')^2 \leq G^2$). Uses `mul_le_mul_of_nonneg_right` and `linarith`. This implies the non-squared part of $g(k) \geq 0$ when $w \geq \delta > 0$.