

Convex Volatility Interpolation

Published in *Risk* magazine, Cutting Edge, February 2026

WBS 22nd Quantitative Finance Conference

Valletta

Fabrice Deschâtres

2 October 2026

VOL*ptima*

- **Follow a model**

- Parsimonious by design: a handful of parameters per expiry against 100+ listed strikes → potentially **under-parameterized**.
 - Parameters enter nonlinearly → **non-convex** calibration.
- We learn from models, but do not apply them directly to vol fitting.

- **Don't follow a model**

- More flexibility, but **nothing does the work for you** → the calibration carries the burden.
- Can use convex optimization (not automatic, but now available), and then afford to be **over-parameterized**.
- **Modular**: encode vol surface properties directly in the calibration.

→ **CVI sits in the second category.**

Today: the *Risk* paper. Some comments on running it in production on slide 35.

Under the radar until recently. Why?

- **Practical frictions**

- Before **Clarabel** (2022), an open-source solver: buy a commercial license, or live with slower, less reliable convergence.
- **Complexity**: the calibration is low-level, properties encoded one at a time.

- **A misconception about splines**

- “Splines oscillate.” They do, unless the fit is regularized and constrained.

- **A gap in the literature**

- QP in **price** space is well covered; **variance** space was not.
- **Strike arbitrage** in the linearly extrapolated wings was left untreated.

Why Variance Space?

- Absence of static arbitrage = no calendar-spread + no strike arbitrage.
 - **No-calendar-spread arbitrage:** linear in both price and variance.
 - **No-butterfly arbitrage:** linear in price, but not in variance.
 - **Price-space drawbacks:**
 - Wing instability (OTM prices $\rightarrow 0$, small absolute errors $\rightarrow$ large vol errors).
 - No intuitive parameters on prices.
- **Variance space is worth the trade-off:**
- Calendar constraints are still linear.
 - Butterfly constraints become nonlinear $\rightarrow$ we linearize them.
 - Intuitive parameters (skew, convexity) are natural in variance space.

Using Basis Functions to Leverage Convex Optimization

- Objective: least squares on the **variance at each listed strike** $\rightarrow$ quadratic in $v(K_j)$.
- Quadratic *in the optimization variables* requires v to be **linear in its parameters**:

$$v(K) = \sum_i w_i \phi_i(K)$$

- This restricts us to a space of **basis functions**.
 - **Prior work in price space**:
 - Fengler (2009): constrained natural cubic spline smoothing of call prices.
 - Lucic (2019): B-splines with bid-ask bounds via QP.
- $\rightarrow$ Similar idea here, but in **variance space**.

Basis Functions

- Choice: **degree 3 B-splines** (the basis functions of cubic splines) in **normalized log-moneyness z** .
- Linear extrapolation in the wings.





Figure: The 9 degree-3 B-spline functions associated with the 7 knots $[-6, -4, -2, 0, 2, 4, 6]$.

Notation

- $k := \log(K/F)$ log-forward moneyness
- $z := \frac{k}{\sigma_\star \sqrt{T}}$ normalized log-moneyness
- $\sigma_\star$ anchor ATM volatility (estimated prior to fitting)
- $v(z, T) := \sigma^2(z, T)$ variance; vT total variance
- Shape function: $z \mapsto \frac{v(z)}{v_\star}$ (variance normalized by $v_\star := \sigma_\star^2$)
- $s := \frac{1}{v_\star} \frac{\partial v}{\partial z}$ dimensionless skew
- $c := \frac{1}{v_\star} \frac{\partial^2 v}{\partial z^2}$ dimensionless convexity
- $d_{1,2} := \frac{-k \pm \frac{vT}{2}}{\sqrt{vT}}$ the Black-Scholes d_1, d_2
- m expiries: $T_1, \dots, T_m$

The Shape Function

- $z \mapsto v(z)/v_*$: variance normalized by anchor ATM variance.
- Approximately the same shape across all expiries.
- Allows comparing smile shapes independent of ATM level and time.




Shape function $v(z)/v_*$ across expiries (S&P 500).

The CVI Cubic Spline

The cubic spline has $n + 2$ parameters per expiry:

- $v(z = 0)$: the ATM variance (1 parameter)
- $s(0) = \frac{1}{v_\star} \frac{\partial v}{\partial z} \Big|_{z=0}$: the ATM skew (1 parameter)
- $\left\{ c(z_i) = \frac{1}{v_\star} \frac{\partial^2 v}{\partial z^2} \Big|_{z_i} \right\}_{0 \leq i \leq n-1}$: the n convexities at knots $z_0, \dots, z_{n-1}$

Edge convexities $c(z_0) = 0$ and $c(z_{n-1}) = 0$ $\rightarrow$ linear extrapolation in tails.



CVI cubic spline (left), first derivative (middle), second derivative (right). Red dots = parameters.

Constructing the Cubic Spline

The $n + 1$ cubic segments have $4(n + 1)$ degrees of freedom, determined by $4(n + 1)$ relations:

Relation	Count
v'' matches convexities at each z_i (left and right)	$2n$
Value matches $v(z=0)$ (left and right of origin)	2
v' matches skew at $z=0$ (left and right of origin)	2
Continuity of v and v' at the other knots ($z_i \neq 0$)	$2(n - 1)$
Cubic coefficient = 0 at edge segments (linear tails)	2
Total	$4(n + 1)$

→ A $4(n + 1) \times 4(n + 1)$ **linear system** maps the $n + 2$ parameters to the polynomial coefficients.

The same variance function can be expressed as:

- **Cubic spline parameters $\mathbf{p}$** (intuitive):
 - ATM variance, ATM skew, convexities at knots
- **B-spline weights $\mathbf{w}$** :
 - $v(z) = \sum_i w_i B_i(z)$ for $z \in [z_0, z_{n-1}]$
 - Linear extrapolation beyond edge knots

Both sets of parameters are linked by a **linear map** $\mathbf{w} = M \mathbf{p}$. σ_* and the knots are fixed before the fit, so M is **known in advance**.

Four penalty terms, each scaling as a **chi-squared statistic**:

- ① **Least Squares Penalty**: fit to the variance implied by the mid price.
 - ② **Above Ask Penalty**: penalize fitting above the ask.
 - ③ **Below Bid Penalty**: penalize fitting below the bid.
 - ④ **Strike Regularization**: smoothness across strikes.
- All expiries fitted **simultaneously**: one global convex optimization.
 - Each penalty is the sum of **dimensionless** numbers across expiries.
 - Consistent behavior regardless of the number of expiries, options per expiry, or volatility level.


Least Squares Penalty

$$\sum_{j=1}^m \frac{1}{N_{T_j, \text{mid}}} \sum_{i \in \mathcal{S}_{T_j, \text{mid}}} (v(K_i, T_j) - v_{\text{mid}}(K_i, T_j))^2 w_{i,j}$$

Weights inversely proportional to the squared bid-ask spread in variance:

$$w_{i,j} = \frac{1}{(v_{\text{ask}}(K_i, T_j) - v_{\text{bid}}(K_i, T_j))^2}$$

→ Exact χ^2 : tight quotes weigh more, wide quotes negligible.



KOSPI 200: wide quotes around ATM do not distort the fit.

Above Ask and Below Bid Penalties

Above Ask: penalizes $\max(v(K_i, T_j) - v_{\text{ask}}(K_i, T_j), 0)^2$

- Weights proportional to vega:

$$w_{i,j,\text{ask}} = q_j \frac{\text{Vega}_{\text{ask}}(K_i, T_j)}{\sum_{l \in \mathcal{S}_{T_j, \text{ask}}} \text{Vega}_{\text{ask}}(K_l, T_j)}$$

- q_j scales the penalty to be consistent with the chi-squared of the least squares term.

Below Bid: symmetric (mirrors Above Ask).

Both terms:

- Capture information from **one-sided quotes** (especially OTM wings).
- Trade-off: require auxiliary variables, increase problem size.
- Practical choice: apply only to strikes with bid or ask, but not both.

Strike Regularization Penalty



Total variation of the convexity curve:

$$\lambda \sum_{j=1}^m \sum_{i=0}^{n-2} |c_{T_j}[i+1] - c_{T_j}[i]|$$

where λ is the **strike regularization factor** and

$$c_{T_j} = (c(z_0, T_j), \dots, c(z_{n-1}, T_j)).$$

- $\lambda = 0.05$: sensible default across underlyings.
- Interpretation: increase χ^2 by 0.05 to reduce total variation by 1.



Tesla, 29 Jan 2020 12:35 ET, 31 Jan expiry (earnings after the close).

- **Linear Extrapolation:**

- $c(z_0) = 0$ and $c(z_{n-1}) = 0$
- Variance linearly extrapolated in log-strike beyond edge knots.

- **Positivity of the Variance inside the Edge Knots:**

- Impose **B-spline coefficients** $\geq 0 \rightarrow v(z) \geq 0$ everywhere.
- Only for the first expiry (subsequent bounded by calendar constraints).

- **Positivity of the Variance in the Tails:**

- $\frac{\partial v}{\partial z}(z_0) \leq 0$ (non-increasing in left tail)
- $\frac{\partial v}{\partial z}(z_{n-1}) \geq 0$ (non-decreasing in right tail)


All constraints above are **linear**.

No-Calendar-Spread-Arbitrage Constraints

Total variance must increase with maturity for a fixed strike-to-forward ratio:

$$v(K_i, T_j) T_j \leq v\left(K_i \frac{F_{T_{j+1}}}{F_{T_j}}, T_{j+1}\right) T_{j+1}$$

- Applied to r strikes: $(m-1)r$ **linear** constraints.
- Total variance curves must not intersect.



Beyond edge knots, variance is linearly extrapolated. Differentiate the calendar constraint w.r.t. $\log K$:

$$\begin{cases} s(z_0, T_j) \sqrt{v_*(T_j) T_j} \geq s(z_0, T_{j+1}) \sqrt{v_*(T_{j+1}) T_{j+1}} \\ s(z_{n-1}, T_j) \sqrt{v_*(T_j) T_j} \leq s(z_{n-1}, T_{j+1}) \sqrt{v_*(T_{j+1}) T_{j+1}} \end{cases}$$

- Equivalently: $T \frac{dv}{dk}$ is non-increasing at z_0 and non-decreasing at z_{n-1} .
- Only $2(m-1)$ additional **linear** constraints.

No-Butterfly-Arbitrage: The Challenge

From Martini and Mingone (2022): $\sigma(k) d_1'(k) d_2'(k) + \sigma''(k) \geq 0$.

Change of variables to normalized log-moneyness z :

$$c \geq \frac{v_\star}{2v} s^2 + \sqrt{v_\star T} s - 2 \left(1 + d_1 \sqrt{\frac{v_\star}{v}} s + d_1 d_2 \frac{v_\star}{4v} s^2 \right)$$

where $d_{1,2} = \frac{-k \pm \frac{vT}{2}}{\sqrt{vT}}$.

- This constraint is **nonlinear and nonconvex** in variance space.
- Cannot be enforced directly in a QP.

Linearization of Butterfly Constraints



Linearize around the reference solution (previous iteration):

$$c \geq \alpha_0 + \alpha_1(s - s_{\text{ref}}) + \alpha_2(v - v_{\text{ref}})$$

$$\begin{cases} \alpha_0 = \frac{v_{\star}}{2v_{\text{ref}}} s_{\text{ref}}^2 + \sqrt{v_{\star} T} s_{\text{ref}} - 2 \left(1 + d_{1,\text{ref}} \sqrt{\frac{v_{\star}}{v_{\text{ref}}}} s_{\text{ref}} + d_{1,\text{ref}} d_{2,\text{ref}} s_{\text{ref}}^2 \frac{v_{\star}}{4v_{\text{ref}}} \right) \\ \alpha_1 = \frac{v_{\star}}{v_{\text{ref}}} s_{\text{ref}} + \sqrt{v_{\star} T} - 2 \left(d_{1,\text{ref}} \sqrt{\frac{v_{\star}}{v_{\text{ref}}}} + d_{1,\text{ref}} d_{2,\text{ref}} s_{\text{ref}} \frac{v_{\star}}{2v_{\text{ref}}} \right) \\ \alpha_2 = \frac{s_{\text{ref}}}{v_{\text{ref}}} \left(-\frac{v_{\star}}{2v_{\text{ref}}} s_{\text{ref}} - 2 \frac{k}{v_{\text{ref}} T} \left(\sqrt{v_{\star} T} - \frac{k}{2} s_{\text{ref}} \frac{v_{\star}}{v_{\text{ref}}} \right) \right) \end{cases}$$

- Now **linear** in $(v, s, c) \rightarrow$ compatible with QP.
- Two-step calibration: (1) without butterfly constraints, (2) with linearized constraints.

Effect of No-Butterfly-Arbitrage Constraints



Without butterfly constraints

Kinks out of the bid-ask; density < 0

AAPL 19 Jan 2024 expiry (6 Dec 2023).

With butterfly constraints

Tracks the bid-ask; density ≥ 0

Lee's Tail Slope Bounds

Lee (2004): $\beta_{R,L} := \limsup_{k \rightarrow \pm\infty} \frac{v(k) T}{|k|} \in [0, 2]$.

Linear wings $\Rightarrow$ the wing slope is constant and *equals* β , so it can be imposed at the knot: $\beta = |T \partial v / \partial \log K|$. We impose the **strict** form $\beta < 2$:

$$s(z_{n-1}) < \frac{2}{\sqrt{v_* T}} \quad s(z_0) > -\frac{2}{\sqrt{v_* T}}$$

Why strict, and not Lee's $\beta \leq 2$:


- $\beta = 2$ is the boundary; strictness stays off it. Strict Lee + linear wings give $d_1 \rightarrow -\infty$ as $k \rightarrow +\infty$ (large-strike no-arb.) and $d_2 \rightarrow +\infty$ as $k \rightarrow -\infty$ (no mass at zero); both collapse to 0 at $\beta = 2$.
 - It pushes $d_1 d_2 \rightarrow +\infty$: past $|k| > \sqrt{v T} \sqrt{1 + \frac{v T}{4}}$ the wing is in the regime $d_1 d_2 > 1$, and the edge knots sit past that.
- $\rightarrow$ That regime is what the **edge-knot butterfly constraint** needs.

At edge knots, $c = 0$. The butterfly condition simplifies to a quadratic in s :

$$v_{\star} \left(\frac{d_1 d_2 - 1}{2v} \right) s^2 + \sqrt{v_{\star}} \left(\frac{2d_1}{\sqrt{v}} - \sqrt{T} \right) s + 2 \geq 0$$

- Discriminant: $\Delta = v_{\star} \left(T + \frac{4}{v} \right) > 0$ (always two real roots).
- Edge knots placed far enough into the tails so that $d_1 d_2 > 1$ (positive quadratic coefficient).

PDF ≥ 0 at Edge Knots: No-Arbitrage Region



Right edge z_{n-1} : $0 < s_- < s_+$. Left edge z_0 is the mirror ($s_- < s_+ < 0$; keep $s \geq s_+$).

- Outside both roots $\Rightarrow$ PDF ≥ 0 , but only the smaller- $|s|$ branch stays **non-empty** under strike-arbitrage-freeness.
- Indeed, every skew on the larger- $|s|$ branch carries a **vertical-spread arbitrage** (CDF outside $[0, 1]$) $\Rightarrow$ empty ($\rightarrow$ next 3 slides).

No Strike Arbitrage: What CVI Enforces

Carr and Madan (2005): a slice of call prices $K \mapsto C(K)$ is free of strike arbitrage iff


	On call prices	Meaning	In CVI
Butterfly	$\partial^2 C / \partial K^2 \geq 0$	PDF ≥ 0	imposed
Vertical spread	$-1 \leq \partial C / \partial K \leq 0$	$0 \leq \text{CDF} \leq 1$	implied
Large strike	$C(K) \rightarrow 0$ as $K \rightarrow \infty$	no mass at ∞	strict Lee
Zero strike	$C(0) = F$	mean = F	automatic

Why vertical spreads come free (Black-Scholes prices; Roper 2010):

- A convex call price that tends to 0 can never increase: $\partial C / \partial K \leq 0$.
- Black-Scholes prices satisfy $(F - K)^+ \leq C \leq F$, so convexity also gives $\partial C / \partial K \geq -1$.

→ So in our setting a vertical-spread arbitrage cannot occur. If we force the fit into a region where one cannot be avoided (say $s \geq s_+$ at the right edge knot, next two slides), no arbitrage-free smile exists (**empty** set of solutions).

The Two Roots Across the Smile flat vol, $T = 1$, $\sigma_\star = 20\%$, $c = 0$



- Dashed s_L^{right} ($d'_1 = 0$) and dotted s_L^{left} ($d'_2 = 0$), Lucic's bounds (next slide), both **strictly between** the roots on each wing: $s_- < s_L^{\text{right}} < s_L^{\text{left}} < s_+$.
- Hatched: no butterfly arbitrage, but past the bounds, so **empty** (next slide).

Why the Larger Branch Is Empty

1. No butterfly arbitrage (Martini–Mingone 2022): $\sigma d'_1 d'_2 + \sigma'' \geq 0$, $' = \partial/\partial k$.
2. No strike arbitrage: d_1 and d_2 decrease along the smile (Fukasawa 2012).

Lucic (2021) writes this as bounds on the skew:

$$s_L^{\text{right/left}} := \frac{4v\sqrt{T}}{\sqrt{v_*}(2k \pm vT)}.$$

3. d'_1 and d'_2 are affine in s and vanish at these bounds:

$$-\sqrt{vT} d'_{1,2} = 1 - s/s_L^{\text{right/left}}.$$

4. At the edge knot $c = 0$, so $\sigma\sigma'' = -\sigma'^2$ and the edge quadratic factors:
 $2(1 - s/s_L^{\text{right}})(1 - s/s_L^{\text{left}}) - \frac{v_*}{2v} s^2 \geq 0$. At either bound one factor vanishes: the quadratic is $-\frac{v_*}{2v} s^2 < 0$ there. It opens upward ($d_1 d_2 > 1$), so

$$s_- < s_L^{\text{right}} < s_L^{\text{left}} < s_+.$$

5. Which vertical-spread arbitrage, by wing:

Left wing ($s_L^{\text{right}} < 0$):

$$\partial P/\partial K \geq 0 \Rightarrow d'_1 < 0 \iff s > s_L^{\text{right}}$$

Larger branch $s \leq s_-$: a put spread.

→ In both wings the larger branch is **empty**.

Right wing ($s_L^{\text{right}} > 0$):

$$\partial C/\partial K \leq 0 \Rightarrow d'_1 < 0 \iff s < s_L^{\text{right}}$$

Larger branch $s \geq s_+$: a call spread.

Argument adapted from V. Lucic.

The Edge Constraint in Closed Form

Only the smaller- $|s|$ root survives. It gives a closed-form skew cap at each edge (needs to be linearized, like the interior constraints):

- At z_{n-1} (right tail): enforce s below the lower root (positive):

$$s \leq \frac{2v\sqrt{T}\left(k - \sqrt{vT} \sqrt{1 + \frac{vT}{4}}\right)}{\sqrt{v_*}\left(k^2 - \frac{v^2T^2}{4} - vT\right)} \quad (\text{i.e. } s \leq s_-)$$

- At z_0 (left tail): enforce s above the upper root (negative):

$$s \geq \frac{2v\sqrt{T}\left(k + \sqrt{vT} \sqrt{1 + \frac{vT}{4}}\right)}{\sqrt{v_*}\left(k^2 - \frac{v^2T^2}{4} - vT\right)} \quad (\text{i.e. } s \geq s_+)$$



→ Only need to enforce at z_0 and z_{n-1} , **not beyond**.

- Given **strict** Lee's tail slope bounds, the RHS reaches a constant asymptotically from below ($k \rightarrow +\infty$) and from above ($k \rightarrow -\infty$).

Next slide: why.

Why Only at Edge Knots?

Past the **right** edge knot the variance is linearly extrapolated, so the skew s is **constant**, while the cap s_- keeps moving.



- **Saturated:** past the knot the numerator of $\partial_z s_-$ keeps the sign it has at the knot, which $d_1 d_2 > 1$ makes positive: the cap rises all the way.
 - **Strictly inside:** the cap *can* fall instead, but never reaches the wing: at a crossing the wing would be saturated *there*, and the same computation sends it back up.
 - Mirror at z_0 : the floor s_+ falls, approached from above.
- The skew stays flat while the cap moves away from it: if the constraint holds at the knot, it holds along the whole wing. **Absence of butterfly arbitrage propagates along a linear wing.**

Putting It All Together: Canonical QP

CVI fits all expiries by solving a QP with linear constraints:

$$\begin{aligned} & \text{minimize} && \frac{1}{2} \mathbf{x}^T P \mathbf{x} + \mathbf{q}^T \mathbf{x} \\ & \text{subject to} && \mathbf{l} \leq A \mathbf{x} \leq \mathbf{u} \end{aligned}$$

Constraint	Linearized	Discretized
Linear Extrapolation	–	–
Positivity of the Variance	–	✓
Positivity in the Tails	–	–
No Calendar Spread	–	✓
No Calendar Spread (Tails)	–	–
PDF ≥ 0 (interior)	✓	✓
Lee's Tail Slope Bounds	–	–
PDF ≥ 0 (edge knots)	✓	–

Two-step calibration handles the linearized constraints.


Numerical Results

	$n_{\text{vol quotes}}$	n_{expiries}	$n_{\text{constraints}}$
SX5E	3 153	19	948
SPX	14 538	46	2 298

Asset	n_{knots}	P Non-Zeros	A Non-Zeros	Time (s)
SX5E	10	1 046	14 424	0.034
	20	1 766	22 465	0.061
	40	5 836	51 710	0.155
SPX	10	2 905	32 158	0.078
	20	6 706	52 683	0.152
	40	30 717	122 416	0.467

- Solver: **Clarabel**, default tolerance; time totalled across 2 calibration iterations.
- SPX: ~ 0.15 s with 20 knots $\times$ 46 expiries (920 parameters, 14 538 quotes).
- Scaling approximately linear in typical use cases.


Convergence of the Fit with the Number of Knots



- The fit converges when the chi-squared statistic plateaus.
- Too few knots acts similarly to applying stronger regularization.
- Mid-to-high teens sufficient for highly liquid assets; 10 for less liquid ones.

Chi-squared (average across expiries for one snapshot) vs. knots per expiry.

Example: S&P 500 during Covid



S&P 500, 16 March 2020 $\rightarrow$ 19 June 2020 expiry ($F \approx 2479$).

Fit and density (top), c as a function of z (bottom).

The spike in c is several times its usual level.

More fits: volptima.com/gallery

CVI: volatility surface fitting framework via convex optimization in variance space.

Summary:

- Arbitrage-free in strike and time, including in the tails.
- All expiries fitted simultaneously.
- Hundredths to tenths of a second calibration.
- Consistent across underlyings via dimensionless parameterization.

Key contributions:

- Intuitive cubic spline parameterization with B-spline dual.
- Chi-squared-scaled objective with bid-ask awareness.
- Derived and linearized butterfly constraints in variance space.
- Treatment of strike arbitrage in the linearly extrapolated wings.

- **What the paper gives you:**

- Fits market data well, no hyperparameter tuning.
- Arbitrage-free.

- **What a production system adds on top:**

- Coupling across expiries, beyond the calendar constraint.
- Wing stabilization and better tail control.
- Memory: a single snapshot is not enough when market makers widen quotes (sell-offs, events, ...).

→ *None of this invalidates the framework (CVI is a low-level calibration encoding a vol surface's properties). Building a production-grade system takes substantial engineering beyond the paper, which is what Volptima's fitter does.*

- CVI paper:
volptima.com/convex-volatility-interpolation.pdf
- More about butterfly arbitrage in variance space / skew quadratic:
volptima.com/maximum-skew-constraint.pdf
- Volptima's CVI fitter:
volptima.com/product
- Contact:
fabrice.deschatres@volptima.com

Q & A