IMPERIAL

Imperial College London

Department of Mathematics

A Finite-Strike Replication Framework for Seasoned and Capped Volatility Swaps

Author Omar BENNANI (CID: 06022447)

A thesis submitted for the degree of

MSc in Mathematics and Finance, 2024-2025

Declaration

The work contained in this thesis is my own work unless otherwise stated.

Acknowledgements

I am deeply grateful to my bank supervisor, **Irina Dolgova**, for her trust, clear guidance, and unwavering focus on what "production-ready" really means. I also thank my school supervisor, **Cristopher Salvi**, for his patient mentorship and insightful questions.

My warmest thanks go to the *Exotics* team for their generous help in developing and validating the model—especially **Hatem Jlassi**, **Ghislain Vong**, who sat with me through every stage of the build and review, and **Alexandre Prevost** from the Structuring team for his pragmatic feedbacks. I am likewise thankful to everyone in the **QIS Investable Indices** team for data support, sanity checks, and many helpful discussions.

Finally, I owe a great deal to my family and friends for their constant encouragement and understanding during this project. Any remaining errors are mine alone.

Abstract

We develop a front-to-back framework to *value and hedge seasoned, capped volatility swaps* using only vanilla options. The methodology is *underlying-agnostic* and applies to both *equity indices and single-name stocks.* On the replication side, we derive a Carr–Madan–type representation for the seasoned payoff and show that a volatility cap induces a transparent boundary term. Continuous integrals are converted into *finite-strike* baskets via a split variance-strike grid aligned with the law of future realised variance; fixings use business days, while discounting/time-to-maturity use calendar-day year fractions.

For calibration and scenarios, we model the future realised variance X as lognormal. Uncapped maturities admit a closed-form calibration; capped maturities are obtained by solving a two-equation system for (μ, ν) (mean and vol-of-variance). Each maturity is pinned by exactly *two* market points (variance- and volatility-swap fair strikes), making the fitted distribution *effectively skewless.* Discounting follows OIS in the trade's collateral currency.

We document *convergence and stability*: PV errors vanish as the strike grid is refined; production runs use large N, so discrepancies are *not* numerical. Benchmarking against an internal bank pricer shows tight co-movement across regimes. Residual differences stem mainly from (i) construction of capped/uncapped fair strikes and (ii) the benchmark's skewed dynamics versus our two-point (skewless) law. A controlled experiment—overriding our inputs with the benchmark's and re-pricing the benchmark with skew $= 0$—*reduces but does not eliminate* the gap. In all cases, remaining pricing differences are < 0.1 vol, which is operationally acceptable for trading and risk.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Volatility has matured from a latent parameter inferred from option markets into a *tradable asset class* with its own microstructure, benchmarks, and specialist investor base. A large body of work underpins this evolution: state-price extraction and static replication via option strips Breeden and Litzenberger (1978); Carr and Madan (1998), log-contract and model-free variance ideas Neuberger (1994), surface construction and no-arbitrage parametrisations Dupire (1994); Gatheral (2006); Gatheral and Jacquier (2014), and the design of investable volatility indices and swaps Demeterfi *et al.* (1999); Carr and Lee (2008); Carr and Wu (2006). This thesis develops a theory–to–implementation pipeline for valuing and hedging *seasoned, capped* volatility derivatives—specifically volatility swaps whose observation windows are partially realised and whose payoffs may be capped in volatility units. The practical goal is to convert any such quote into a *tradable basket of European options* on a finite strike grid, augmented, when necessary, by a light dynamic overlay, in a way that is accurate on real markets and robust under stress.

1.1 Macro Landscape for Volatility Derivatives

Since the late 1990s, *tradability of volatility* has progressed from implicit exposures (gamma scalping, vega-weighted books) to explicit line-item instruments: **variance swaps**, **volatility swaps**, corridor/conditional variants, and investable indices linked to option strips Demeterfi *et al.* (1999); Carr and Lee (2008); Carr and Wu (2006). By the mid-2020s, open interest spans flagship indices (S&P 500, Euro Stoxx 50, Nikkei 225), regional indices, and a large single-name universe. Advances in surface modelling and extraction Dupire (1994); Gatheral (2006); Gatheral and Jacquier (2014); Fengler (2009) and exchange/e-RFQ plumbing have lowered minimum sizes and improved price discovery, enabling large-scale static replication.

Three structural forces underpin this expansion:

- **Regulation and capital.** Post-crisis adoption of collateralised discounting and multi-curve frameworks made pricing and risk more transparent across desks Piterbarg (2010); Henrard (2014). Mapping exotic payoffs to vanillas improves netting and simplifies risk controls.

- **Buy-side demand.** Institutions use VS/VolS for tail hedging, carry harvesting, dispersion and correlation views; index-linked and single-name exposures coexist Carr and Wu (2006); Carr and Lee (2008).

- **Technology & market plumbing.** Efficient surface smoothing/extrapolation with no static arbitrage Gatheral and Jacquier (2014); Fengler (2009) plus investable definitions ISDA (2011) make replication operational at scale.

Episodes such as *Volmageddon* (2018), the COVID shock (2020), and the 2022–2025 risk-off waves highlight both the opportunity and the challenge: convex payoffs are valuable, but their *nonlinear* and sometimes *path-dependent* nature makes *accurate, operationally simple* replication essential Broadie and Jain (2008); Cont and Tankov (2004).

1.2 Instrument Taxonomy: Variance vs. Volatility Swaps

Variance and volatility swaps allow investors to trade future realised fluctuation against a fixed strike; they differ in payoff *linearity* and thus in replication.

Variance swaps (VS) settle on the annualised *average of squared log returns.* Under standard assumptions, an exact static replication exists using a continuum of OTM options—the log-contract/BL machinery and Carr–Madan integral Breeden and Litzenberger (1978); Neuberger (1994); Carr and Madan (1998); Demeterfi *et al.* (1999). FFT and related transforms make these computations efficient on grids Carr and Madan (1999).

Volatility swaps (VolS) settle on the *square root* of realised variance. Concavity induces a *convexity adjustment* between the vol-swap strike and the square root of the variance-swap strike; replication mixes a static strip with a small dynamic overlay or a model-based convexity factor Demeterfi *et al.* (1999); Broadie and Jain (2008); Carr and Lee (2008).

Seasoning and caps. Two structuring features are ubiquitous in practice:

- **Seasoning:** Trades are often entered mid-window. The remaining payoff references only the *future* path. Day-count, business-day sampling and fixing conventions are codified in the ISDA Equity Variance Swap Definitions and must be respected in pricing ISDA (2011).
- **Caps:** Many solutions cap the payout in *volatility* units to limit tail risk. On the replication side this introduces boundary terms in Carr–Madan–type representations and interacts with sparse wings on real surfaces Demeterfi *et al.* (1999); Lee (2004); Gatheral (2006).

1.3 Desk Mandate and Project Genesis

In June 2025 I joined the **Quantitative Investment Strategies (QIS) Investable Indices** desk at Bank of America. A significant share of the desk's vega notional involves *seasoned and capped* VS/VolS—often embedded in corporate programmes, capital-efficiency notes, and dispersion/correlation structures mixing index VS shorts and single-name VolS longs. Trading, risk, and XVA functions require a *production-ready* engine that (i) ingests a seasoned/capped quote and the live option surface, (ii) produces a *finite-strike* static replication with weights, (iii) quantifies convexity/jump/model risk, and (iv) runs intraday.

1.4 Research Gap and Academic Motivation

The literature provides elegant continuous-strike identities and model insights Breeden and Litzenberger (1978); Neuberger (1994); Carr and Madan (1998); Demeterfi *et al.* (1999); Carr and Lee (2008), but desk-grade replication of *seasoned, capped* instruments faces practical gaps:

- **Seasoning.** Partial realisation changes the payoff to $f(aX+b)$ of the *future* realised variance X, with a, b determined by *fixing-day counts*—a detail mandated by market definitions ISDA (2011) yet often simplified in academic treatments.
- **Caps.** A volatility cap maps to a *variance cutoff* that adds a boundary term in replication; tractable, finite-strike versions compatible with market wings are rarely spelled out Demeterfi *et al.* (1999); Lee (2004); Gatheral (2006).
- **Sparse wings and finite grids.** Real surfaces vanish beyond $\approx \pm 3\sigma$. Continuous integrals must be evaluated on finite grids without static arbitrage; SVI/no-arbitrage smoothing and targeted quadrature help but require robust recipes Gatheral and Jacquier (2014); Fengler (2009); Lee (2025).
- **Jump and convexity risk.** Square-root (VolS) and min (caps) amplify jump sensitivity. Diffusion-only convexity corrections can misprice stressed scenarios unless tail risk is modelled explicitly Broadie and Jain (2008); Cont and Tankov (2004).

Addressing these gaps advances both *practice* (replication accuracy, risk control) and *theory* (seasoned/capped functional replication on finite grids).

1.5 Thesis Objectives

The thesis delivers a front-to-back solution with two integrated objectives:

(A) Analytical & Replication Framework Derive a Carr–Madan–type representation for seasoned and capped VolSwap payoffs, exhibiting the $f(x) = \sqrt{ax+b}$ structure from seasoning and an explicit cap boundary term. Map continuous integrals to *finite* strike baskets with weights consistent with no-arbitrage surfaces Carr and Madan (1998); Gatheral and Jacquier (2014).

(B) Calibration & Testing Adopt a parsimonious *lognormal* market model for the *future realised variance* X, pinned by variance/volatility swap quotes (uncapped: closed form; capped: two-equation system), consistent with empirical vol-of-variance dynamics Bühler (2022) and operational under collateralised discounting Piterbarg (2010); Henrard (2014).

1.6 Scope, Assumptions, and Practical Constraints

The replication layer is *model-agnostic*, requiring only arbitrage-free vanilla premia. For calibration and scenario analysis we use a lognormal law for X (positive, right-skew, tractable moments), which is easily stressed via vol-of-var and jump overlays Carr and Lee (2008); Bühler (2022); Cont and Tankov (2004). Key working assumptions:

- Business-day fixings for the past–future split; calendar-day year fractions for discounting/time-to-maturity ISDA (2011); Piterbarg (2010); Henrard (2014).
- Surface liquidity concentrated near-the-money; wings smoothed/extrapolated with no static arbitrage Gatheral and Jacquier (2014); Fengler (2009); Lee (2004).
- Replication on discrete strikes with controlled truncation/quad error and stable weights Carr and Madan (1999); Lee (2025).

1.7 Roadmap

The thesis proceeds from theory to implementation and validation.

Chapter 2 — Theoretical Framework. We formalise instrument conventions and seasoning on business-day grids (§2.2), proving the decomposition $\mathrm{RV}(0,T) = aX + b$ (Lemma 2.2.1). We specify a lognormal market law for the *future* realised variance X and its calibration (§2.3): uncapped maturities admit a closed-form link between $(K_{\mathrm{var}}, K_{\mathrm{vol}})$ and the vol-of-variance ν, while capped maturities are recovered by a two-equation system ((2.3.6)–(2.3.7)). We derive option premia on X (Proposition 2.3.1) and obtain Carr–Madan–type functional replication for the seasoned payoff $f(x) = \sqrt{ax+b}$ (Theorem 2.4.1, §2.4). For capped contracts, we exhibit an explicit *cap adjustment* boundary term (Proposition 2.5.1). Finally, we convert the continuous integrals into a *finite-strike* implementation (§2.6) using a split z-grid (puts left, calls right), variance-strike mapping, and truncation rules consistent with the law of X.

Chapter 3 — Data, Implementation Tests, and Empirics. We describe the EURO STOXX 50 dataset and swap quotes (capped/uncapped) used for calibration and pricing (§3.1). Numerical robustness is established via grid-refinement experiments (§3.2), showing PV convergence as strike density increases. We handle discounting under CSA-consistent OIS curves and document its behaviour over time (§3.3). We then analyse price dynamics across maturities and strikes, including stress episodes (§3.4). A comprehensive benchmark vs. the bank's internal pricer follows (§3.5), together with a *same-data / no-skew* experiment isolating sources of residual discrepancy. We conclude the empirical section by comparing our fitted vol-of-variance to the benchmark's (§3.6).

Appendices. Technical proofs (the Black integral for calls and the Carr–Madan representation) are provided for completeness, alongside a full instrument registry used in the experiments (tickers, maturities, caps, and conventions).

1.8 Contributions

This thesis closes the gap between continuous–strike theory and desk–grade replication of *seasoned, capped* volatility swaps. The main contributions are:

- **Finite–strike grid.** An explicit strike grid is designed for Carr–Madan replication: split at the variance forward, mapped via a normal z–grid, and truncated with universal tail rules.
- **Cap adjustment.** A volatility cap is handled by stopping the integral at the variance cutoff and adding a single boundary term (*CapAdj*), making cap risk explicit and hedgeable.
- **Seasoning consistency.** The framework embeds business–day fixing counts and calendar–day discounting, aligning pricing with ISDA conventions.
- **Validation.** Numerical tests confirm convergence; empirical benchmarking against a bank pricer shows residual spreads below 0.1 vol, acceptable for trading and risk.

Chapter 2

Theoretical Framework

2.1 Volatility–Swap Payoff

A volatility swap with notional M and strike K_{vol} pays at maturity

$$\mathrm{Payoff}_{\mathrm{VolS}} = M\big(\mathrm{RVOL}(0,T) - K_{\mathrm{vol}}\big),$$

with an optional cap implemented as $M\big(\min\{\mathrm{RVOL}(0,T), \sigma_{\cap}\} - K_{\mathrm{vol}}\big)$.
$\mathrm{RVOL}(0,T)$ denotes the annualised *realised volatility* over the observation window $[t_0, T]$.

2.2 Set–up, Seasoning and Notation

Let t_0 be the *Inception Day*, t the current *Calculation Day*, and T the *Maturity Day.*

Denote by $\mathrm{CD}_{\mathrm{past}}(t_0, t)$ and $\mathrm{CD}_{\mathrm{total}}$ the numbers of *business* fixing days in $[t_0, t]$ and $[t_0, T]$ respectively.

We also use $\mathrm{CalDays}(t, T)$ for the number of *calendar* days from t to T (weekends and holidays included).

2.2.1 Elapsed-time weight and residual tenor

$$\begin{aligned} w_t &:= \frac{\mathrm{CD}_{\mathrm{past}}(t_0, t)}{\mathrm{CD}_{\mathrm{total}}}, \\ \tau &:= \frac{\mathrm{CalDays}(t, T)}{365}. \end{aligned} \tag{2.2.1}$$

Interpretation

- **Fixing-day weighting.** Volatility contracts fix on *business* days; the past–future split must therefore use the count of fixings. Thus $w_t \in [0, 1]$ is exactly the fraction already realised.
- **Financial time.** Pricing/discounting runs in *years*; using $\tau = \mathrm{CalDays}(t, T)/365$ aligns with day-count conventions and the later parametrisation $\sigma_X = \nu\sqrt{\tau}$. In short, w_t weights the realised data; τ sets the horizon for forward pricing and convexity.

2.2.2 Past realised volatility

$$\mathrm{RVOL}_{\mathrm{past}}(t_0, t) = \sqrt{\frac{252}{\mathrm{CD}_{\mathrm{past}}(t_0, t)} \sum_{i=1}^{\mathrm{CD}_{\mathrm{past}}(t_0,t)} r_i^2}\,. \tag{2.2.2}$$

Here $r_i := \ln(P_i/P_{i-1})$ are close–to–close *log* returns on business (fixing) days, and P_i is the closing price on the i-th fixing day. We also use variance units:

$$\mathrm{RV}_{\mathrm{past}} \;:=\; \mathrm{RVOL}^2_{\mathrm{past}}. \tag{2.2.3}$$

Interpretation

- **Log returns.** $\ln(P_i/P_{i-1})$ are time-additive, matching the Carr–Madan replication framework and avoiding the compounding bias of simple returns.
- **Annualisation (252).** Under a stationary daily-volatility regime with $\mathbb{E}[r_i] = 0$ and $\mathrm{Var}(r_i) = \sigma^2_{\mathrm{day}}$, the trading-year variance is $\sigma^2_{\mathrm{ann}} \approx 252\,\sigma^2_{\mathrm{day}}$. The factor $252/\mathrm{CD}_{\mathrm{past}}$ rescales the sample mean of squared daily returns to an annualised variance.

Consistency properties. By construction $w_t = 0$ at inception and $w_t = 1$ at maturity; consequently $\tau = \mathrm{CalDays}(t, T)/365$ decreases to 0 as we approach T. These limits will imply:
(i) in the unseasoned case ($w_t = 0$) all randomness lies in the future realised variance;
(ii) at $t = T$, $\mathrm{RV}(0, T) = \mathrm{RV}_{\mathrm{past}}$ and the model's randomness vanishes as $\tau \to 0$.

Lemma 2.2.1 (Seasoning decomposition)**.** *Let* $X := \mathrm{RV}(t, T)$ *be the* future realised variance *over the residual window* $[t, T]$. *Then the full-period realised variance satisfies*

$$\mathrm{RV}(0, T) \;=\; a\,X + b, \qquad a := 1 - w_t, \quad b := w_t\,\mathrm{RV}_{\mathit{past}}. \tag{2.2.4}$$

Proof. The full-period average of squared returns on $[t_0, T]$ splits into the weighted average of the past part and the future part, with weights equal to the fractions of fixing days. □

2.3 Market Model for the Future Realised Variance

2.3.1 Model and parameterisation

We model the (strictly positive) future realised variance X as lognormal with

$$\ln X \sim \mathcal{N}\big(m,\; v\big), \qquad v := \nu^2\tau, \qquad m := \ln\mu - \tfrac{1}{2}v, \tag{2.3.1}$$

where $\mu > 0$ is the target mean, $\nu \geq 0$ is a *log–volatility per* $\sqrt{\text{year}}$, and $\tau > 0$ is the time to maturity (in years).

Implications. With this parametrisation one has

$$\mathbb{E}[X] = e^{m+\frac{1}{2}v} = \mu, \qquad \mathrm{Var}(X) = \mu^2\big(e^v - 1\big).$$

More generally, for any $a \in \mathbb{R}$,

$$\mathbb{E}\big[X^a\big] = \exp\Big(am + \tfrac{1}{2}a^2 v\Big) = \mu^a\,\exp\Big(\tfrac{1}{2}(a^2 - a)\,v\Big). \tag{2.3.2}$$

In particular,

$$\mathbb{E}[\sqrt{X}] = \mu^{1/2}\,\exp\Big(-\tfrac{1}{8}v\Big) = \sqrt{\mu}\,\exp\Big(-\tfrac{1}{8}\nu^2\tau\Big). \tag{2.3.3}$$

Equation (2.3.2) follows by writing $X = e^{m+\sqrt{v}\,Z}$ with $Z \sim \mathcal{N}(0, 1)$ and using $\mathbb{E}[e^{tZ}] = e^{t^2/2}$. Taking $a = \frac{1}{2}$ yields (2.3.3).

Market inputs (capped vs. uncapped). For each maturity and pricing day, we use the pair of fair strikes matching the instrument's cap:

$$\big(K_{\mathrm{var}}^{\mathrm{mkt}}, K_{\mathrm{vol}}^{\mathrm{mkt}}\big) = \begin{cases} \big(K_{\mathrm{var}}^{\mathrm{unc,\,mkt}}, K_{\mathrm{vol}}^{\mathrm{unc,\,mkt}}\big), & \text{uncapped,} \\ \big(K_{\mathrm{var}}^{\mathrm{cap,\,mkt}}(C), K_{\mathrm{vol}}^{\mathrm{cap,\,mkt}}(C)\big), & \text{capped at } C. \end{cases}$$

2.3.2 Calibration

Uncapped case (closed form). Set $\mu = K_{\mathrm{var}}^{\mathrm{unc,\,mkt}}$ and $K_{\mathrm{vol}}^{\mathrm{unc}} = K_{\mathrm{vol}}^{\mathrm{unc,\,mkt}}$. From (2.3.3),

$$\nu^2 \;=\; \frac{8}{\tau}\,\ln\!\left(\frac{\mu}{\big(K_{\mathrm{vol}}^{\mathrm{unc}}\big)^2}\right)$$

Capped case (two–equation system in μ and ν). We solve directly for (μ, ν), where $\mu := \mathbb{E}[X] = K_{\mathrm{var}}^{\mathrm{unc}}$ and $\ln X \sim \mathcal{N}(m, \nu^2\tau)$ with $m = \ln\mu - \frac{1}{2}\nu^2\tau$.

Step 1 — Write capped strikes as truncated expectations.

$$K_{\mathrm{var}}^{\mathrm{cap}}(C) = \mathbb{E}[\min(X, C^2)], \qquad K_{\mathrm{vol}}^{\mathrm{cap}}(C) = \mathbb{E}[\min(\sqrt{X}, C)].$$

Step 2 — Use the identity $\min(Y, K) = Y - (Y - K)^+$.

$$K_{\mathrm{var}}^{\mathrm{cap}}(C) = \mathbb{E}[X] - \mathbb{E}[(X - C^2)^+] = \mu - \mathbb{E}[(X - C^2)^+], \tag{2.3.4}$$

$$K_{\mathrm{vol}}^{\mathrm{cap}}(C) = \mathbb{E}[\sqrt{X}] - \mathbb{E}[(\sqrt{X} - C)^+]. \tag{2.3.5}$$

Step 3 — Compute $\mathbb{E}[\sqrt{X}]$. Since $\ln X \sim \mathcal{N}(m, \nu^2\tau)$,

$$\mathbb{E}[\sqrt{X}] = \exp\!\Big(\tfrac{1}{2}m + \tfrac{1}{2}\cdot\tfrac{\nu^2\tau}{4}\Big) = \sqrt{\mu}\,\exp\!\Big(-\tfrac{1}{8}\nu^2\tau\Big) =: F_{\mathrm{vol}}.$$

Step 4 — Recall the Black formula for a lognormal call. For a lognormal underlying with forward F, strike K, log–vol σ, maturity τ,

$$\mathrm{Black}(F, K, \sigma, \tau) := F\,\Phi(d_1) - K\,\Phi(d_0), \quad d_1 = \frac{\ln(F/K) + \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}},\; d_0 = d_1 - \sigma\sqrt{\tau},$$

and $\mathrm{Black}(F, K, \sigma, \tau) = \mathbb{E}[(Y - K)^+]$ when Y has $\mathbb{E}[Y] = F$ and $\ln Y$-variance $\sigma^2\tau$.

Step 5 — Map our payoffs to Black.

- *Variance:* $Y = X$ is lognormal with mean μ and log–vol ν, so $\mathbb{E}[(X - C^2)^+] = \mathrm{Black}(\mu, C^2, \nu, \tau)$.
- *Volatility:* $Y = \sqrt{X}$ is lognormal with mean F_{vol} from (2.3.2) and log–vol $\nu/2$[1], hence $\mathbb{E}[(\sqrt{X} - C)^+] = \mathrm{Black}(F_{\mathrm{vol}}, C, \nu/2, \tau)$.

System to solve in (μ, ν). Substitute the Black expressions into (2.3.4)–(2.3.5):

$$K_{\mathrm{var}}^{\mathrm{cap}}(C) = \mu \;-\; \mathrm{Black}\big(\mu,\, C^2,\, \nu,\, \tau\big)\,, \tag{2.3.6}$$

$$K_{\mathrm{vol}}^{\mathrm{cap}}(C) = \underbrace{\sqrt{\mu}\,\exp\!\Big(-\tfrac{1}{8}\nu^2\tau\Big)}_{F_{\mathrm{vol}}} \;-\; \mathrm{Black}\Big(\sqrt{\mu}\,e^{-\nu^2\tau/8},\, C,\, \tfrac{\nu}{2},\, \tau\Big)\,. \tag{2.3.7}$$

Equations (2.3.6)–(2.3.7) define a two–equation system in (μ, ν) (with $\mu = \mathbb{E}[X]$). We solve it numerically (e.g., a 2D solver, or a 1D root find in ν and back out μ); the call terms are increasing in ν, ensuring a unique solution in practice.

[1] Since $\ln X \sim \mathcal{N}(m, \nu^2\tau)$ and $Y = \sqrt{X} = e^{\frac{1}{2}\ln X}$, we have $\ln Y = \frac{1}{2}\ln X \sim \mathcal{N}(m/2,\ (\nu^2\tau)/4)$. Hence the log–variance scales by 1/4 and the log–vol by 1/2, i.e., $\nu/2$. Consequently $\mathbb{E}[Y] = \exp\big(m/2 + \frac{1}{8}\nu^2\tau\big) = \sqrt{\mu}\,e^{-\nu^2\tau/8}$, consistent with (2.3.2).

2.3.3 Options on realised variance

Proposition 2.3.1 (Undiscounted call/put on X)**.** *Let X be as in* (2.3.1)*, with* $\ln X \sim \mathcal{N}(m, \nu^2\tau)$ *and* $F_{var} := \mathbb{E}[X] = e^{m+\frac{1}{2}\nu^2\tau}$*. Then the undiscounted European premia are*

$$\mathrm{Call}(k) = \mathbb{E}[(X-k)^+] \;=\; \mathrm{Black}(F_{var},\, k,\, \nu,\, \tau)\,, \tag{2.3.8}$$

$$\mathrm{Put}(k) = \mathbb{E}[(k-X)^+] \;=\; \mathrm{Black}(F_{var},\, k,\, \nu,\, \tau) \;-\; \big(F_{var} - k\big), \tag{2.3.9}$$

where $\mathrm{Black}(F, K, \sigma, \tau) := F\,\Phi(d_1) - K\,\Phi(d_0)$ *is the forward Black call, and the put follows by (undiscounted) put–call parity. (Discounted prices multiply both premia by the discount factor.)*

2.4 Carr–Madan Functional Replication

Theorem 2.4.1 (Carr–Madan representation Carr and Madan (1998))**.** *Let $X \geq 0$ be integrable with $F := \mathbb{E}[X] < \infty$. Let $f : (0,\infty) \to \mathbb{R}$ be C^2 such that f and f' have at most linear growth and the integrals $\int_0^F |f''(k)|\,\mathrm{Put}(k)\,dk$ and $\int_F^\infty |f''(k)|\,\mathrm{Call}(k)\,dk$ are finite, where* $\mathrm{Put}(k) := \mathbb{E}[(k-X)^+]$ *and* $\mathrm{Call}(k) := \mathbb{E}[(X-k)^+]$ *are* undiscounted *option premia. Then*

$$\mathbb{E}[f(X)] \;=\; f(F) + \int_0^F f''(k)\,\mathrm{Put}(k)\,dk + \int_F^\infty f''(k)\,\mathrm{Call}(k)\,dk. \tag{2.4.1}$$

Uncapped seasoned volatility swap. Apply Theorem 2.4.1 to the concave $f(x) := \sqrt{ax+b}$ from (2.2.4), after checking the hypotheses:

(i) $X = \mathrm{RV}(t,T) \geq 0$ is the future realised variance, hence non–negative and integrable; and $F_{\mathrm{var}} = \mathbb{E}[X] < \infty$.

(ii) $f(x) = \sqrt{ax+b}$ with $a = 1 - w_t \in (0,1]$ and $b = w_t\,\mathrm{RV}_{\mathrm{past}} \geq 0$ is C^2 on $(0,\infty)$,

$$f'(x) = \frac{a}{2\sqrt{ax+b}} > 0, \qquad f''(x) = -\,\frac{a^2}{4(ax+b)^{3/2}}\,,$$

and has at most linear growth ($f(x) \sim \sqrt{a}\,x^{1/2}$ as $x \to \infty$).

(iii) The Carr–Madan integrability conditions hold: for $k \downarrow 0$, $\mathrm{Put}(k) \leq k$ and $|f''(k)| \leq \frac{a^2}{4(ak+b)^{3/2}}$, so $|f''(k)|\,\mathrm{Put}(k) \lesssim k^{-1/2}$ if $b = 0$ (integrable) and is bounded if $b > 0$. For $k \to \infty$, $\mathrm{Call}(k) \leq F_{\mathrm{var}}$ and $|f''(k)| \lesssim k^{-3/2}$, hence $\int_{F_{\mathrm{var}}}^\infty |f''(k)|\,\mathrm{Call}(k)\,dk < \infty$.

With $h(x) := -f''(x) = \dfrac{a^2}{4(ax+b)^{3/2}}$, Theorem 2.4.1 yields

$$\mathbb{E}\big[\sqrt{aX+b}\big] \;=\; f(F_{\mathrm{var}}) - \int_0^{F_{\mathrm{var}}} h(k)\,\mathrm{Put}(k)\,dk - \int_{F_{\mathrm{var}}}^\infty h(k)\,\mathrm{Call}(k)\,dk. \tag{2.4.2}$$

Multiply by the discount factor gives the present value.

2.5 Capped Volatility Swap and the Cap Adjustment

Let $\sigma_\cap$ be the cap in volatility units and define the corresponding variance cutoff $k_\cap$ by $f(k_\cap) = \sigma_\cap$, i.e.

$$k_\cap = \frac{\sigma_\cap^2 - b}{a}. \tag{2.5.1}$$

We write the capped payoff as

$$\min\{f(X), \sigma_\cap\} = f(X) - (f(X) - \sigma_\cap)^+.$$

For any increasing C^1 function g and any K, we have [2]

$$(g(X) - g(K))^+ = \int_K^\infty g'(k)\, \mathbf{1}_{\{X>k\}}\, dk.$$

Taking expectations and using $\partial_k \mathrm{Call}(k) = -\mathbb{P}(X > k)$[3],

$$\mathbb{E}\big[(g(X) - g(K))^+\big] \;=\; \int_K^\infty g'(k)\,\big(-\partial_k \mathrm{Call}(k)\big)\, dk \;\overset{4}{=}\; g'(K)\,\mathrm{Call}(K) + \int_K^\infty g''(k)\,\mathrm{Call}(k)\, dk.$$

With $g = f$ and $K = k_\cap$,

$$\mathbb{E}\big[(f(X) - \sigma_\cap)^+\big] \;=\; f'(k_\cap)\,\mathrm{Call}(k_\cap) + \int_{k_\cap}^\infty f''(k)\,\mathrm{Call}(k)\, dk.$$

Subtracting this from (2.4.2) yields the following.

Proposition 2.5.1 (Capped seasoned vol-swap)**.**

$$\mathbb{E}\big[\min\{f(X), \sigma_\cap\}\big] = f(F_{var}) - \int_0^{F_{var}} h(k)\,\mathrm{Put}(k)\, dk - \int_{F_{var}}^{k_\cap} h(k)\,\mathrm{Call}(k)\, dk - \underbrace{f'(k_\cap)\,\mathrm{Call}(k_\cap)}_{CapAdj}. \tag{2.5.2}$$

with $h = -f''$ *and* $f'(x) = \frac{a}{2\sqrt{ax+b}}$.

Proof. Combine (2.4.2) with the identity for $\mathbb{E}[(f(X) - \sigma_\cap)^+]$ above; the tail integral $\int_{k_\cap}^\infty$ cancels and the boundary term $f'(k_\cap)\,\mathrm{Call}(k_\cap)$ remains. □

2.6 Finite–Strike Implementation

The integrals in (2.4.2)–(2.5.2) are over $k \in (0, \infty)$, but we evaluate them on a *finite* strike set that mirrors the law of X and controls truncation error.

Change of variables and grid design

Write $X = F_{\text{var}} \exp\big(\sigma_X Z - \frac{1}{2}\sigma_X^2\big)$ with $Z \sim \mathcal{N}(0,1)$ and $\sigma_X = \nu\sqrt{\tau}$. The map

$$k \;=\; F_{\text{var}} \exp\Big(\nu\sqrt{\tau}\, z - \tfrac{1}{2}\nu^2\tau\Big)$$

sends standard-normal quantiles z to *variance strikes* k. Choosing nodes that are *evenly spaced in* z therefore allocates more strikes where the risk-neutral density of X is largest (near the body) and fewer in the wings, which is near-optimal for integrals that can be written as expectations under the X–law (Carr–Madan).

[2]For a fixed realisation x, since g is increasing and C^1, $(g(x) - g(K))^+ = \mathbf{1}_{\{x>K\}} \int_K^x g'(k)\, dk$ by the fundamental theorem of calculus. Rewrite the last integral as $\int_K^\infty g'(k)\, \mathbf{1}_{\{x>k\}}\, dk$. Replacing x by the random variable X yields the stated identity.

[3]Since $\mathrm{Call}(k) = \mathbb{E}[(X-k)^+]$, we have $\frac{\partial}{\partial k}(x-k)^+ = -\mathbf{1}_{\{x>k\}}$ for fixed x. Interchanging derivative and expectation (e.g., by Dominated Convergence) gives $\partial_k\,\mathrm{Call}(k) = -\,\mathbb{E}[\mathbf{1}_{\{X>k\}}] = -\,\mathbb{P}(X > k)$.

[4]Integration by parts with $u = g'(k)$, $dv = -\partial_k \mathrm{Call}(k)\, dk$, so $v = \mathrm{Call}(k)$, $du = g''(k)\, dk$. The boundary term at ∞ vanishes under mild growth of g' and $\mathrm{Call}(k) \to 0$, leaving $g'(K)\mathrm{Call}(K) + \int_K^\infty g''(k)\mathrm{Call}(k)\, dk$.

Split at the forward. The Carr–Madan representation naturally splits at $k = F_{\rm var}$ into a put integral for $k < F_{\rm var}$ and a call integral for $k > F_{\rm var}$. This avoids type switching inside any panel and makes both integrands smooth and monotone in their respective domains, improving trapezoidal accuracy.

Lower and upper bounds in z

We take

$$z_i \in [z_{\min}, z_{\max}], \qquad z_{\min} = -5, \qquad z_{\max} = \begin{cases} +5, & \text{uncapped case,} \\ z_\cap, & \text{capped case,} \end{cases}$$

with

$$z_\cap = \frac{\ln(k_\cap/F_{\rm var}) + \frac{1}{2}\nu^2\tau}{\nu\sqrt{\tau}}.$$

Why $z_{\min} = -5$? Under $Z \sim \mathcal{N}(0,1)$, the left-tail probability $\Phi(-5) \approx 2.87\times10^{-7}$. The put-panel integrand is $h(k)\text{Put}(k)$ with $h(k) = \frac{a^2}{4(ak+b)^{3/2}}$ decreasing in k and $\text{Put}(k) \le k$. Hence the *lower-tail truncation error* is bounded by

$$\int_0^{k_{\min}} h(k)\,\text{Put}(k)\,dk \;\le\; \int_0^{k_{\min}} \frac{a^2}{4(ak+b)^{3/2}}\,k\,dk \;=\; \frac{1}{2}\Big(\sqrt{ak_{\min}+b} - \sqrt{b}\Big),$$

where $k_{\min} = F_{\rm var} e^{-5\nu\sqrt{\tau}-\frac{1}{2}\nu^2\tau}$. For typical parameters this is $< 10^{-6}$ of notional. A distributional bound gives the same conclusion: since $k \le k_{\min}$ corresponds to $Z \le -5$, $\text{Put}(k) \le k_{\min}\,\mathbb{P}(Z \le -5)$ and $\int_0^{k_{\min}} h(k)\text{Put}(k)dk \lesssim h(0)\,k_{\min}\,\Phi(-5)$.

Why $z_{\max} = +5$ in the *uncapped* case? For $k > F_{\rm var}$ the integrand is $h(k)\text{Call}(k)$. Using $\text{Call}(k) \le \mathbb{E}[X] = F_{\rm var}$ and $h(k) \le \frac{a^{1/2}}{4\,k^{3/2}}$ for k large, the *upper-tail truncation error* beyond $k_{\max}$ satisfies

$$\int_{k_{\max}}^{\infty} h(k)\,\text{Call}(k)\,dk \;\le\; \frac{a^{1/2}F_{\rm var}}{4}\int_{k_{\max}}^{\infty} k^{-3/2}\,dk \;=\; \frac{a^{1/2}F_{\rm var}}{2\sqrt{k_{\max}}}.$$

With $k_{\max} = F_{\rm var}e^{5\nu\sqrt{\tau}-\frac{1}{2}\nu^2\tau}$ (i.e. $z_{\max} = +5$), this bound is of order $\frac{a^{1/2}}{2}e^{-\frac{1}{2}(5\nu\sqrt{\tau}-\frac{1}{2}\nu^2\tau)}$, which is $\ll 10^{-5}$ for the equity parameter range ($\nu\sqrt{\tau} \lesssim 1.5$). Thus $z_{\max} = +5$ makes the truncation error negligible compared to bid–ask.

Why $z_{\max} = z_\cap$ in the *capped* case? When a volatility cap $\sigma_\cap$ is present the variance cutoff is $k_\cap = (\sigma_\cap^2 - b)/a$; the Carr–Madan integrals stop at $k_\cap$ and the remaining tail is accounted for exactly by the boundary term $f'(k_\cap)\,\text{Call}(k_\cap)$ (the CapAdj). Hence there is *no* truncation on the right: we set the upper bound to the image of $k_\cap$, namely $z_\cap$.

Split z-grid with $2N{+}1$ nodes

To align with the put/call split of the Carr–Madan integrals and avoid double-counting at $k = F_{\rm var}$, we use an asymmetric grid in z:

- **Left (put) panel:** $z \in [z_{\min}, 0]$ with $N{+}1$ *nodes* including 0.

$$\Delta z_- \;=\; \frac{0 - z_{\min}}{N}, \qquad z_i^- \;=\; z_{\min} + i\,\Delta z_-, \quad i = 0,1,\dots,N \;\;(\Rightarrow\; z_N^- = 0).$$

- **Right (call) panel:** $z \in (0, z_{\max}]$ with N *nodes* excluding 0.

$$\Delta z_+ \;=\; \frac{z_{\max} - 0}{N}, \qquad z_j^+ \;=\; j\,\Delta z_+, \quad j = 1, 2, \dots, N \;\; (\Rightarrow\; z_1^+ > 0).$$

  Here

$$z_{\max} \;=\; \begin{cases} +5, & \text{uncapped case,} \\ z_\cap = \dfrac{\ln(k_\cap / F_{\mathrm{var}}) + \frac{1}{2}\nu^2\tau}{\nu\sqrt{\tau}}, & \text{capped case.} \end{cases}$$

Mapping nodes to variance strikes uses the transform

$$k \;=\; F_{\mathrm{var}}\,\exp\!\Big(\nu\sqrt{\tau}\,z - \tfrac{1}{2}\nu^2\tau\Big),$$

so that

$$k_i^- \;=\; F_{\mathrm{var}}\exp\!\Big(\nu\sqrt{\tau}\,z_i^- - \tfrac{1}{2}\nu^2\tau\Big), \quad i = 0, \dots, N \quad (\text{put panel, includes } k_N^- = F_{\mathrm{var}}),$$

$$k_j^+ \;=\; F_{\mathrm{var}}\exp\!\Big(\nu\sqrt{\tau}\,z_j^+ - \tfrac{1}{2}\nu^2\tau\Big), \quad j = 1, \dots, N \quad (\text{call panel, strictly } k_j^+ > F_{\mathrm{var}}).$$

Trapezoidal sums with $2N{+}1$ nodes. We approximate each panel separately (monotone integrands, better accuracy):

$$\int_0^{F_{\mathrm{var}}} h(k)\,\mathrm{Put}(k)\,dk \;\approx\; \sum_{i=0}^{N} h\big(k_i^-\big)\,\mathrm{Put}\big(k_i^-\big)\,\Delta k_i^-, \qquad \Delta k_i^- := k_{i+1}^- - k_i^-,$$

$$\int_{F_{\mathrm{var}}}^{k_{\max}} h(k)\,\mathrm{Call}(k)\,dk \;\approx\; \sum_{j=1}^{N} h\big(k_j^+\big)\,\mathrm{Call}\big(k_j^+\big)\,\Delta k_j^+, \qquad \Delta k_j^+ := k_{j+1}^+ - k_j^+,$$

with $k_{\max} = F_{\mathrm{var}} e^{5\nu\sqrt{\tau} - \frac{1}{2}\nu^2\tau}$ in the uncapped case and $k_{\max} = k_\cap$ in the capped case (the remaining tail is then handled exactly by the CapAdj term $f'(k_\cap)\,\mathrm{Call}(k_\cap)$).

Consistency Checks

- **Unseasoned limit.** If $w_t = 0$ then $a = 1$, $b = 0$ and $f(x) = \sqrt{x}$; (2.4.2) reduces to the standard vol-swap Carr–Madan formula with convexity factor $e^{-\nu^2\tau/8}$.
- **No-cap limit.** As $k_\cap \to \infty$, the capped formula (2.5.2) converges to the uncapped formula (2.4.2).
- **Deterministic limit.** As $\nu \to 0$ (hence $\sigma_X \to 0$), option premia vanish and $\mathbb{E}[\sqrt{aX + b}] \to \sqrt{aF_{\mathrm{var}} + b}$ by continuity—this is also the Jensen upper bound.

Chapter 3

Data, Implementation Tests, and Empirics

3.1 Data Description

3.1.1 Universe, sample, frequency

We work on the EURO STOXX 50 (SX5E) and the dealer market for *variance* and *volatility* swaps written on SX5E. The dataset contains daily end-of-day (EOD) index levels and EOD quotes for VolSwaps and VarSwaps across standard maturities. Unless otherwise stated, *capped* quotes follow the desk convention

$$\mathrm{Cap} \;=\; 2.5 \times K_{\mathrm{vol}},$$

i.e., a volatility cap $C = 2.5\,K_{\mathrm{vol}}$ (variance cutoff $k_{\cap} = C^2$). Quotes are treated as undiscounted forwards. Fixings use business days, while time-to-maturity and discounting use calendar-day year fractions, aligning pricing with settlement and the replication layer in Section 2.6.

3.1.2 EURO STOXX 50 spot

The SX5E spot exhibits the familiar large moves: the COVID crash (March 2020), the 2022 energy/inflation shock, and several risk-off waves in 2023–2024. A particularly sharp episode appears in April 2025, widely linked by market commentary to *Trump's tariffs*. This shock drove a brief but pronounced jump in option-implied volatilities and, as we will see, in swap strikes.

3.1.3 Raw quotes and strike fields

Table 3.1 shows an excerpt of the raw EOD panel used for calibration and replication. Capped and uncapped datasets share this schema; every strike used later can be traced back to a specific EOD record with maturity and underlying level.

3.1.4 Volatility-swap strikes over time

Capped and uncapped vol-swap strikes move together in quiet regimes, with the capped series sitting below the uncapped series. During stress, the gap widens because the cap truncates the right tail of realised volatility. The April 2025 tariff headlines generate a clear spike; the cap dampens the peak while preserving its timing.

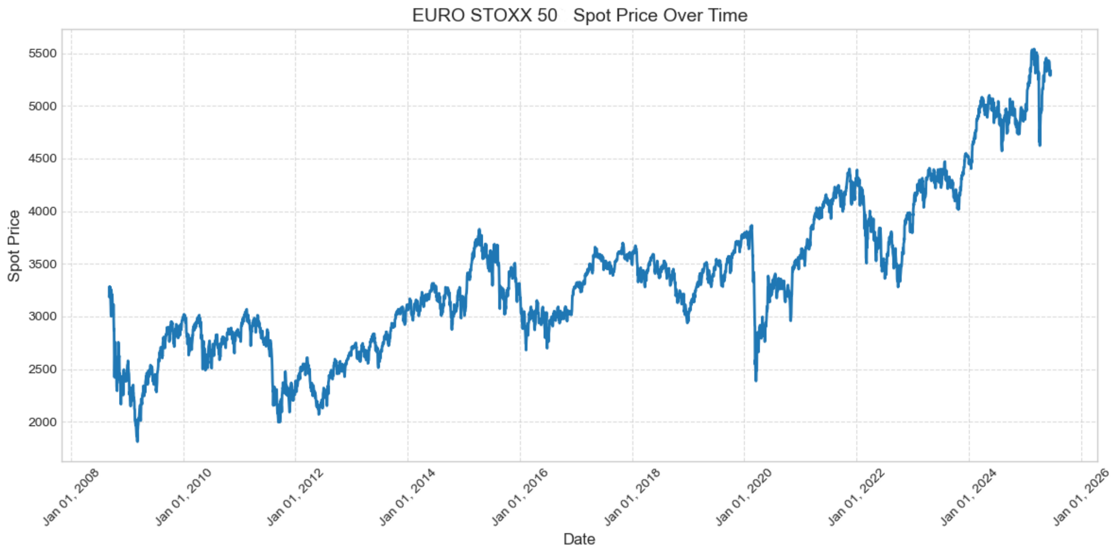

Figure 3.1: EURO STOXX 50 spot index over time.

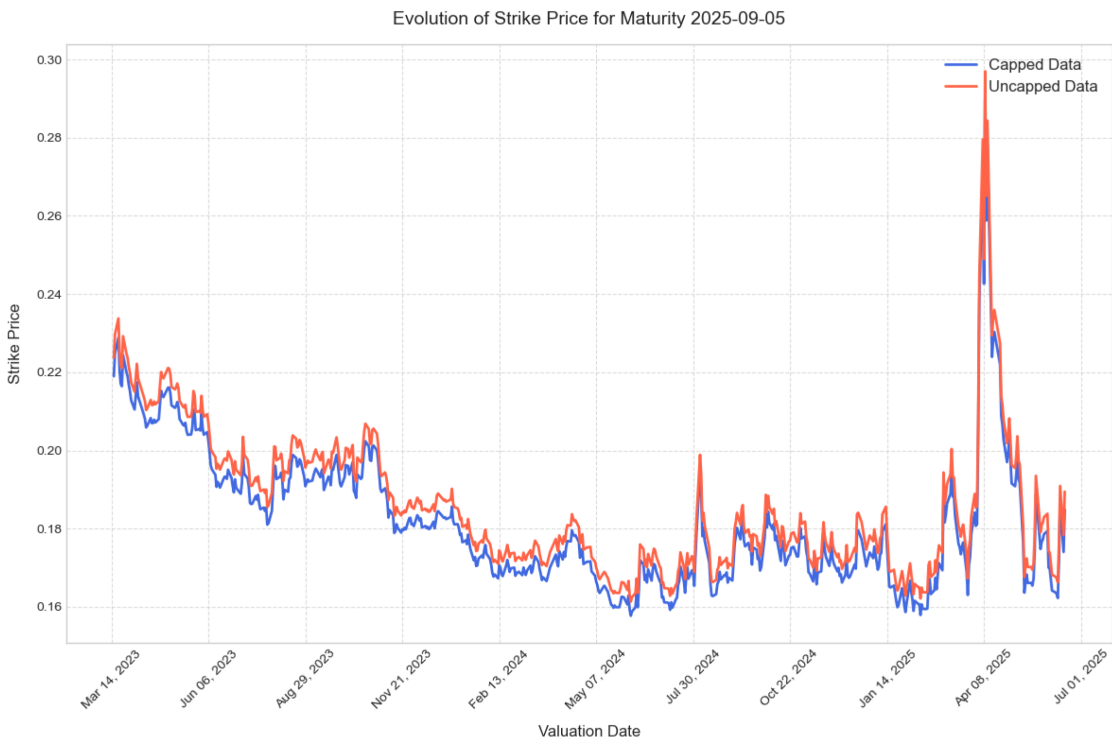

Figure 3.2: Volatility-swap fair strikes (capped vs. uncapped) over time for a representative maturity.

Table 3.1: Calibration panel — sample of EOD quotes for SX5E (excerpt of the raw dataset).

Valuation Date	Market Type	Underlying	Spot Price	Expiration Date	Volatility
09/16/2008	VarSwap	STOXX .50E	3088.430	01/31/2011	0.270
09/12/2008	VolSwap	STOXX .50E	3278.020	05/17/2010	0.252
09/11/2008	VolSwap	STOXX .50E	3222.100	01/11/2010	0.242
09/08/2008	VolSwap	STOXX .50E	3284.120	10/23/2009	0.250
09/16/2008	VarSwap	STOXX .50E	3088.430	06/17/2009	0.270
09/26/2008	VolSwap	STOXX .50E	3156.460	06/08/2010	0.258
09/17/2008	VarSwap	STOXX .50E	3018.770	03/09/2009	0.286
09/16/2008	VarSwap	STOXX .50E	3088.430	05/28/2010	0.253
09/22/2008	VarSwap	STOXX .50E	3181.770	06/23/2009	0.248
09/10/2008	VolSwap	STOXX .50E	3242.020	09/28/2010	0.240
09/15/2008	VarSwap	STOXX .50E	3151.170	08/16/2010	0.259
09/22/2008	VarSwap	STOXX .50E	3181.770	04/28/2009	0.253
09/19/2008	VarSwap	STOXX .50E	3253.520	06/25/2010	0.255
09/08/2008	VolSwap	STOXX .50E	3284.120	04/23/2009	0.239
09/15/2008	VarSwap	STOXX .50E	3151.170	07/05/2010	0.258

Note: This is an excerpt for layout; the full panel contains ∼5.37M rows with the same schema.

3.1.5 Variance-swap strikes over time

For comparability, the variance-swap fair strikes are plotted in *volatility units* (i.e., $\sqrt{K_{\text{var}}}$). On this scale, the April 2025 spike is *just as large* as in the vol-swap series, and the capped/uncapped wedge is of similar magnitude. The economic difference remains: VS payoffs are linear in realised variance, while VolS payoffs are concave; expressing VS in volatility units helps visual comparison but does not alter the underlying payoff convexity.

In the next sections we use these time series to calibrate the lognormal model for future realised variance (Section 2.3) and to test finite-strike replication accuracy.

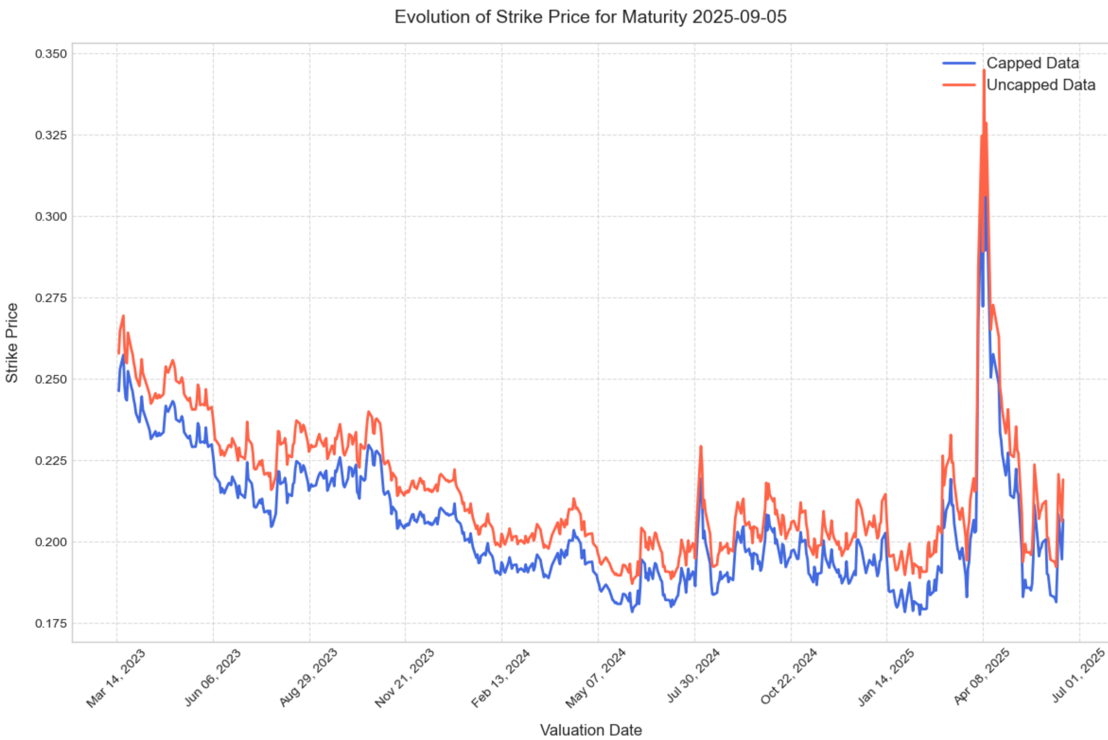

Figure 3.3: Variance-swap fair strikes, expressed in volatility units ($\sqrt{K_{\text{var}}}$), capped vs. uncapped, for a representative maturity.

3.2 Convergence and Stability Testing

3.2.1 Set–up and error metric

To verify numerical robustness, we scale the discretisation of the replication grid by a factor $\alpha \in \{1, 2, \dots, 20\}$. Concretely, the number of option nodes (and the quadrature nodes tied to them) is set to $N(\alpha) = \alpha\, N_0$, so α acts as a uniform refinement of the strike grid used by the Carr–Madan integrals in Section 2.6. We take $\alpha^\star = 40$ as the reference and report the *relative error* for a generic output PV as

$$\text{Error}_\alpha(PV) \;=\; \frac{PV_\alpha - PV_{\alpha^\star}}{PV_{\alpha^\star}}\,. \tag{3.2.1}$$

Unless noted otherwise, figures show percentage values of $\text{Error}_\alpha(\text{PV})$.

3.2.2 PV convergence across instruments

Figures 3.6 display the relative PV error as α increases, for two representative instruments. one instrument is *uncapped* , one is *capped.*

Reading the panels.

- **Uncapped instruments (a).** With a very coarse grid ($\alpha = 1-2$) the relative error is large and negative (undershooting), reflecting insufficient wing coverage and truncation bias in the call/put integrals. As the grid is refined the error increases monotonically toward 0 and becomes essentially flat around $\alpha \gtrsim 8$ (within plotting resolution of 0%). This is the expected *from-below* convergence when integrals are truncated on a too-short grid.
- **Capped instrument (b).** The cap introduces a kink in the payoff and the CapAdj boundary term. At small α we observe an initial overshoot (most visible at $\alpha = 2$), followed by small alternating corrections as the grid captures the cutoff strike more precisely. From $\alpha \approx 5$ the error oscillates tightly around 0 (within a few tenths of a percent), and the curve flattens further as α grows.
- **Rate of convergence.** Across all four cases, PV errors are negligible for $\alpha \geq 10$; in practice $\alpha \in [8, 12]$ provides an excellent compute–accuracy trade-off for intraday use.

3.3 Discounting and Collateral

3.3.1 Collateral–aware discounting

Pricing is performed under OIS discounting consistent with the trade's collateral agreement. Given the collateral currency c, the discount factor applied to any payoff at maturity T observed on date t is

$$D_c(t,T) = \exp\Big(-\int_t^T r_c^{\text{OIS}}(u)\,du\Big), \tag{3.3.1}$$

where r_c^{OIS} is the overnight OIS curve of the collateral currency (EUR: €STR; USD: SOFR). Throughout, market quotes (option premia and swap strikes) are handled as *undiscounted forwards*; the replication value is computed in forward units and multiplied by $D_c(t,T)$ at the end, which keeps discounting consistent and transparent.

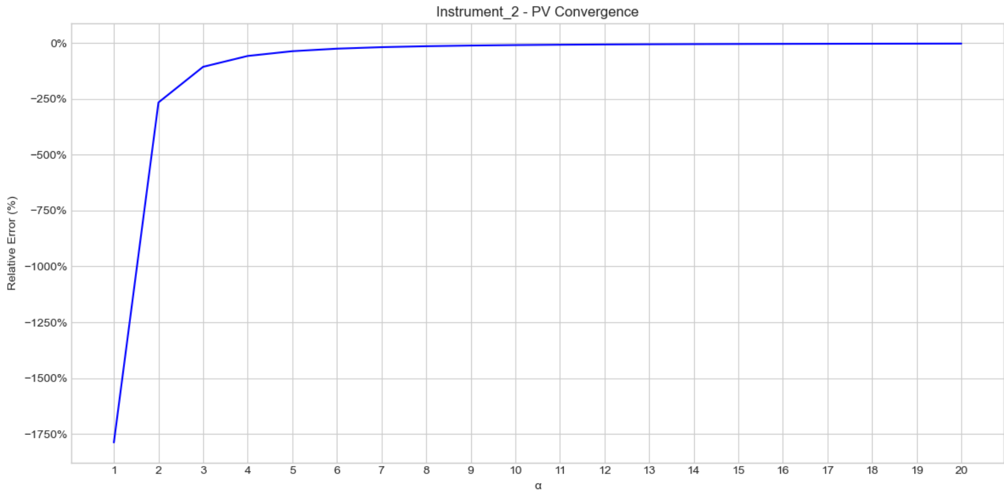

Figure 3.4: *
(a) Instrument 2 — **uncapped**.

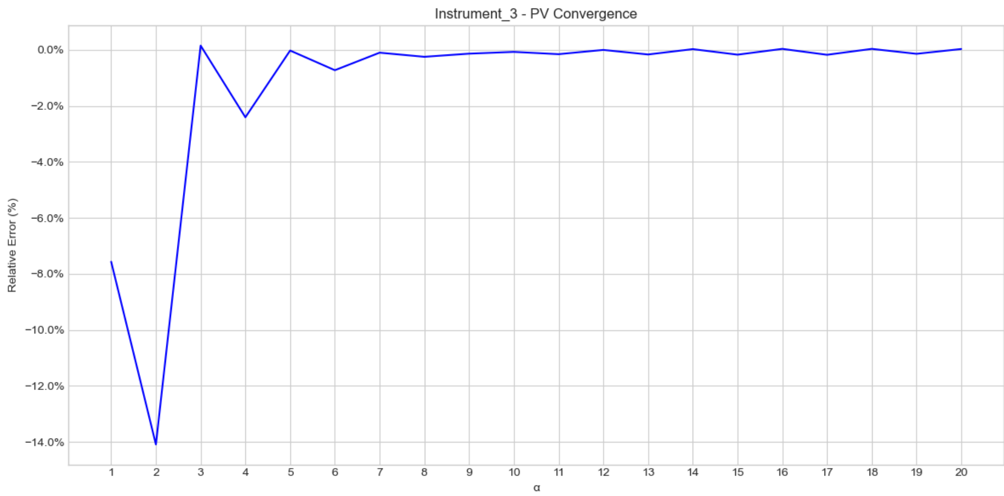

Figure 3.5: *
(b) Instrument 3 — **capped**.

Figure 3.6: PV convergence vs. grid scale α (row 1).

3.3.2 Empirical behaviour

For a fixed instrument, as the valuation date advances toward its maturity, the time to payment shrinks and the discount factor drifts up toward 1. The panel below show this for a representative maturity; the small day-to-day wiggles reflect modest shifts in the OIS curves.

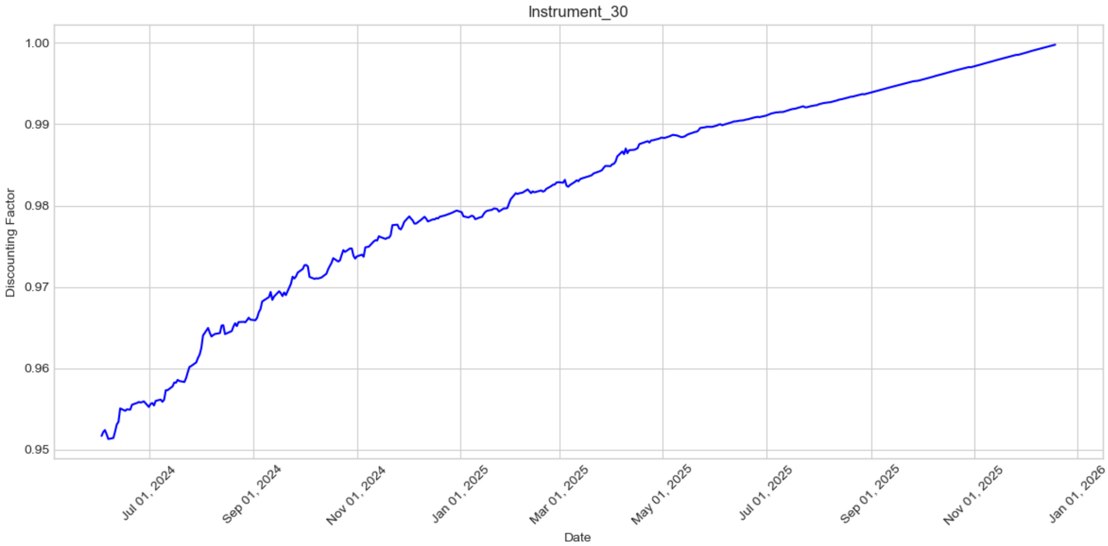

Figure 3.7: Instrument 30 - Discount factor derived from the collateral OIS curves.

The serie start around 0.95 when the residual tenor is long and rise smoothly to 1.00 as maturity approaches. The slope reflects the level of the collateral OIS curve: flatter when overnight rates are stable, steeper when rates fall. Because discounting is applied as a single multiplicative factor in our pipeline, its contribution to replication error is negligible relative to strike-grid discretisation (Section 3.2); nevertheless, CSA-consistent discounting is crucial for correct PVs across instruments booked under different collateral currencies.

3.4 Evolution of Prices Across Maturities

Figures 3.8–3.11 show time–series of model PVs for sets of volatility–swap instruments that *share a maturity* but have *different fixed strikes* K_{vol}. Across panels, the curves move together and exhibit a clear jump in April 2025 (trade–tariff headlines), followed by a decline whose speed depends on the contract and the date window.

Maturity profile. The magnitude and decay of the April move differ across maturities. Some shorter time to maturity revert quickly after the spike, while longer ones remain elevated for longer; in other cases the peak itself is larger at an intermediate/longer tenor. This variation is consistent with the fact that PV combines both the expected future volatility over the residual window and the fixed strike, and that the convexity term in §2.3 depends on τ.

Strike dependence (within a maturity).

- **Near–ATM.** Series plotted with strikes close to the prevailing forward tend to track one another most closely. This is where the option surface is densest, so the finite–strike integrals in §2.6 are most informative.

- **Low vs. high K.** Holding all else equal, higher strikes shift PV downward and lower strikes upward ($\mathrm{PV} \approx \mathbb{E}[\mathrm{RVOL}] - K$). In the figures the ordering by K generally reflects this, although transient deviations can appear around stress days when wings and caps (if any) start to matter more.
- **Near–zero PV regime.** For lines that sit close to zero most of the time (typically high K), small absolute differences can look large in relative terms; brief widening of the dispersion around April is therefore not surprising.

Capped vs. uncapped. In the capped panels (Figures 3.8–3.9) some series appear locally flattened at the peak, which is consistent with the cap becoming relevant near the shock. Outside such episodes—or when the cap is far from realised levels—capped and uncapped series are visually very close, as expected when the boundary term in §2.5 is inactive.

Summary. Overall, the panels are in line with the replication/calibration framework: co-movement across strikes at a given maturity, sensitivity to the April 2025 shock, and maturity/strike patterns that match the role of τ and K. Where dispersion widens, the timing aligns with stressed conditions where wing quotes and cap mechanics carry more weight.

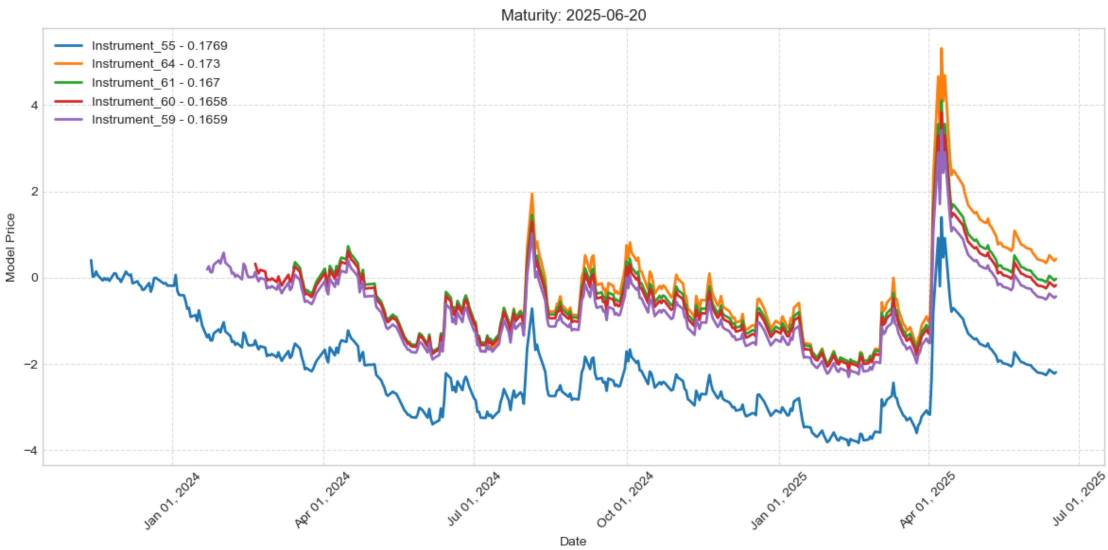

Figure 3.8: Capped VolSwap prices, maturity 2025-06-20, multiple strikes.

3.5 Benchmarking vs. Bank of America's Internal Model

We compare our finite–strike replication against *Bank of America's internal benchmark pricer* for volatility swaps. Both engines are fed the same discounting setup (CSA–consistent OIS; cf. §3.3) and we report daily present values (PVs).

Error metric. For date t, instrument i,

$$\mathrm{AE}_t^{(i)} \;=\; \left|\mathrm{PV}_{t,i}^{\mathrm{model}} - \mathrm{PV}_{t,i}^{\mathrm{bench}}\right|$$

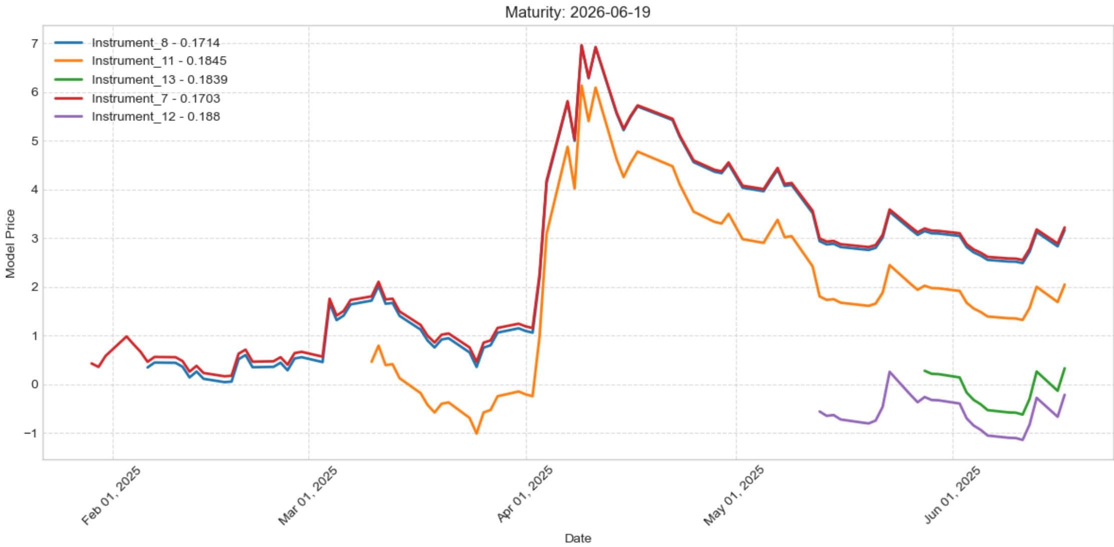

Figure 3.9: Capped VolSwap prices, maturity 2025-12-19, multiple strikes.

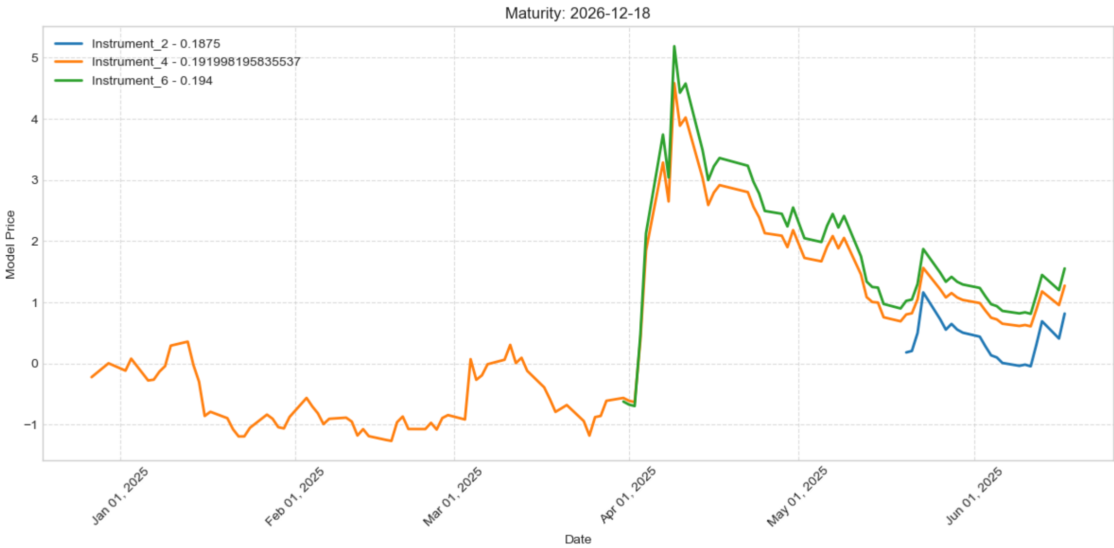

Figure 3.10: Uncapped VolSwap prices, maturity 2025-12-19, multiple strikes.

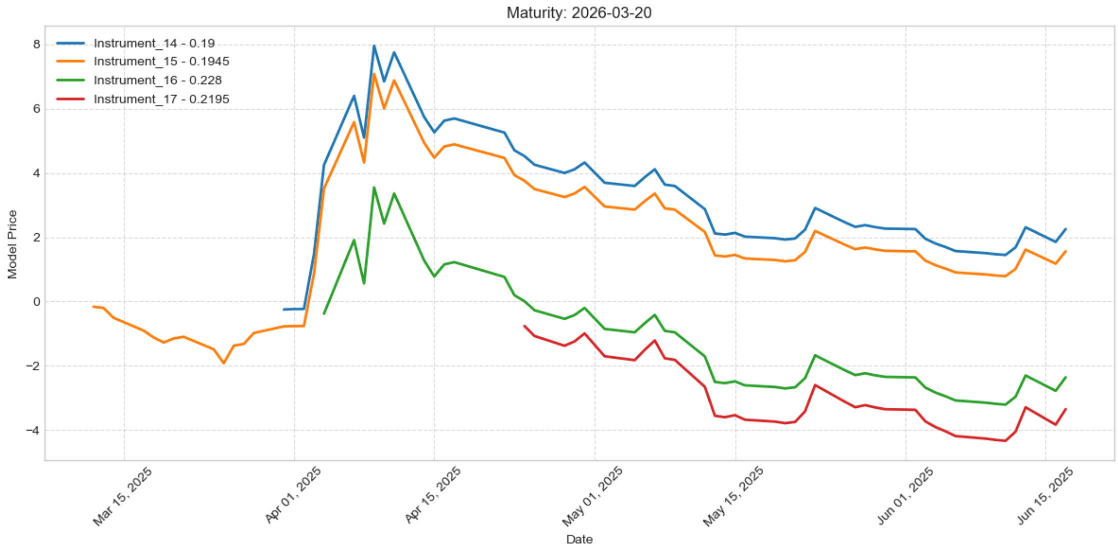

Figure 3.11: Uncapped VolSwap prices, maturity 2026-03-20, multiple strikes.

is plotted together with the two PV curves. This shows both co–movement and the residual mismatch.

Uncapped example (Instrument 26)

Figure 3.14 overlays model and benchmark PVs (top) and the absolute error (bottom) for an *uncapped* VolSwap. Co–movement is tight over the full window; transient widening appears around the April 2025 volatility shock.

Capped example (Instrument 55)

Figure 3.17 repeats the exercise for a *capped* VolSwap. Differences increase precisely when the cap is near binding during the April 2025 spike and shrink as the surface normalises.

Attribution of residual differences (not numerical). The discrepancies are *not* due to discretisation or convergence: prices are computed with a large strike/quadrature grid N, making numerical error negligible. Instead, they arise from:

- **Market inputs (capped vs. uncapped fair strikes).** Small differences in how capped/uncapped fair strikes are built (cap level, day-count, smoothing of EOD quotes) translate directly into PV shifts. When the cap is near binding, these differences are mechanically amplified through the boundary term in the capped payoff.

- **Two–point calibration $\Rightarrow$ effectively skewless model.** Our calibration uses exactly *two* market points per maturity—the variance–swap and volatility–swap fair strikes (capped or uncapped)—to pin the lognormal law of future realised variance X. Concretely, the system $\left(K^{\bullet}_{\text{var}},\, K^{\bullet}_{\text{vol}}\right)$ determines only $\mu = \mathbb{E}[X]$ and the log–volatility ν, with no extra degree of freedom to fit skew/asymmetry of X's distribution. This makes our calibrated law *effectively skewless* relative to the benchmark, which incorporates skewed dynamics; around volatility shocks this structural difference shows up as the modest, directionally consistent gaps seen in the figures.

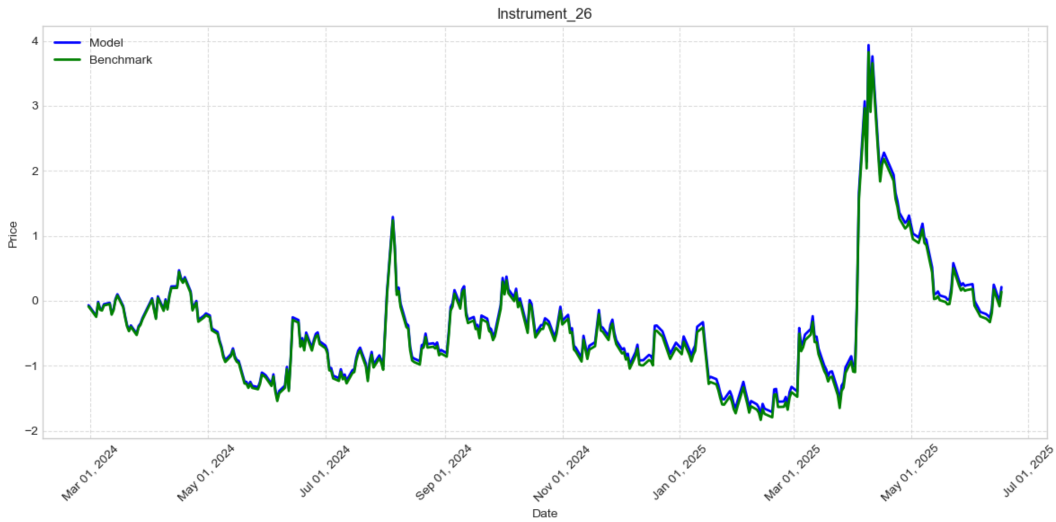

Figure 3.12: *
(a) Prices: model (blue) vs. benchmark (green).

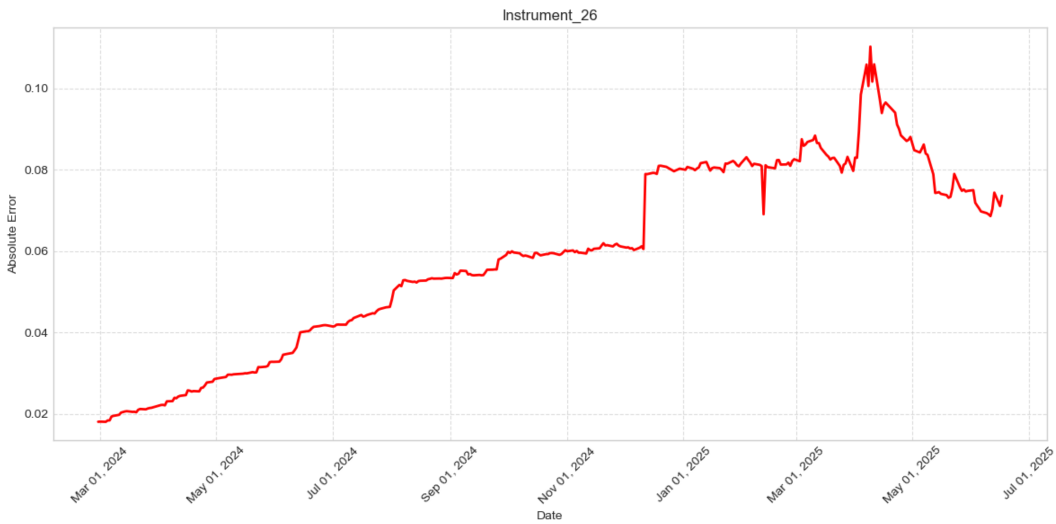

Figure 3.13: *
(b) Absolute error over time.

Figure 3.14: Instrument 26 — **uncapped** VolSwap: PV comparison and absolute error.

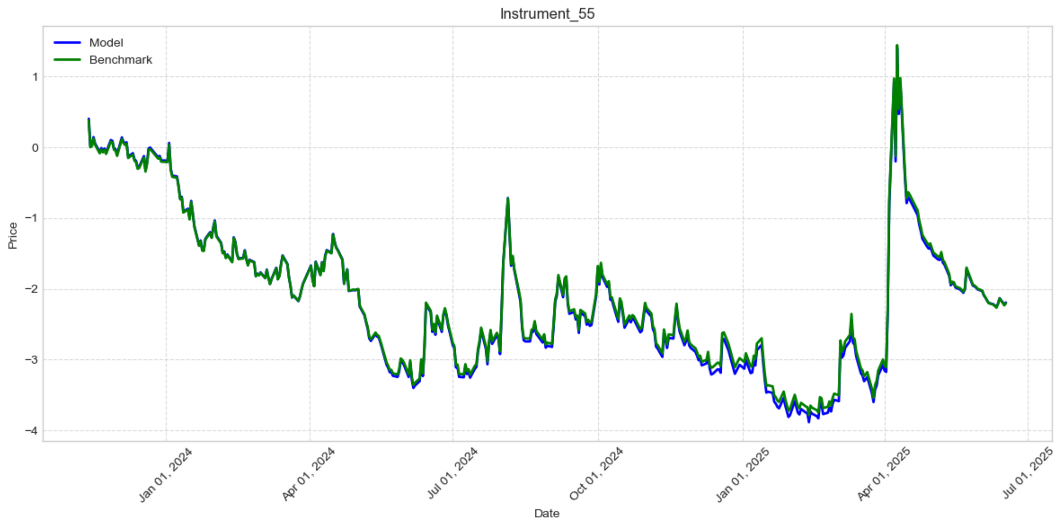

Figure 3.15: *
(a) Prices: model (blue) vs. benchmark (green).

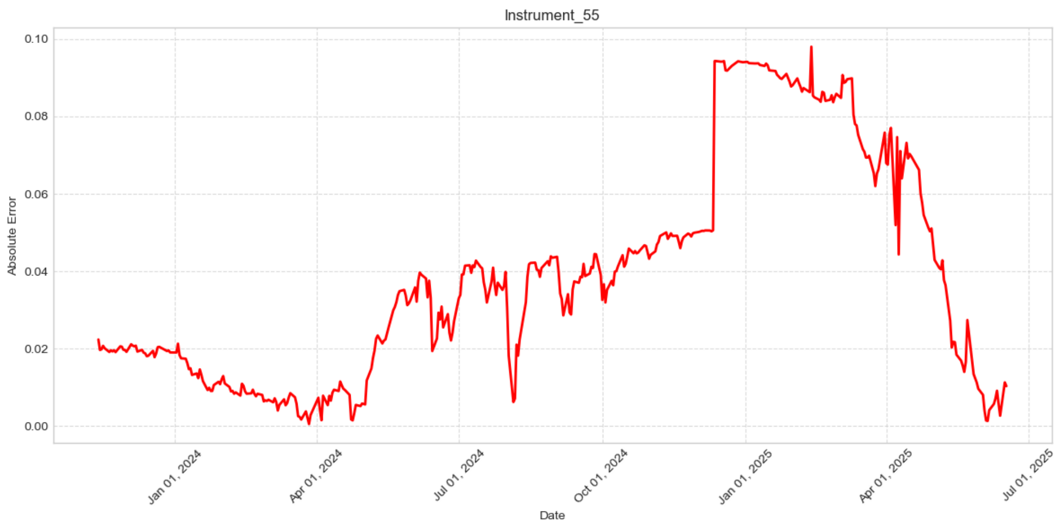

Figure 3.16: *
(b) Absolute error over time.

Figure 3.17: Instrument 55 — **capped** VolSwap: PV comparison and absolute error.

Overall, the fit is strong across regimes; residual spreads are economically small, arise when cap mechanics and distributional skew matter most, and are consistent with the distinct market–data construction and modelling assumptions noted above.

Same–data / no–skew experiment

To isolate the sources of the residual spread, we run a controlled experiment in which we *remove market–data differences* and *neutralise benchmark skew*. Specifically, for each date t and maturity T:

- We build a **model*** price, $\mathrm{PV}_t^{\mathrm{model}^\star}$, by *overriding* our inputs with the benchmark's uncapped variance fair strike and variance–of–variance (i.e., we set $K_{\mathrm{var}}^{\mathrm{unc}}$ and the vol-of-var equal to the benchmark's), keeping our replication and two–point calibration otherwise unchanged.[1]
- We build a **benchmark*** price, $\mathrm{PV}_t^{\mathrm{bench}^\star}$, by re–pricing on the benchmark engine with *skew set to zero* ("skew= 0") while keeping its own market data.

We then compare absolute discrepancies

$$\mathrm{Diff}_t \;=\; \left|\mathrm{PV}_t^{\mathrm{model}} - \mathrm{PV}_t^{\mathrm{bench}}\right|, \qquad \mathrm{Diff}_t^\star \;=\; \left|\mathrm{PV}_t^{\mathrm{model}^\star} - \mathrm{PV}_t^{\mathrm{bench}^\star}\right|.$$

Figures 3.20–3.23 show that $\mathrm{Diff}^\star$ is *generally smaller* than Diff across time, confirming that (i) differences in market–data construction for capped/uncapped fair strikes and (ii) the presence of skew in the benchmark are the main drivers of the baseline spread. Prices are computed with a *large* strike/quadrature grid N, so convergence error is negligible.

Takeaways.

- **Reduction but not elimination.** Using benchmark fair strikes/vol-of-var in model* and shutting skew in bench* reduces the gap materially, but some residual differences remain.
- **Why residuals persist.** Remaining spreads are consistent with structural choices that still differ across engines: cap boundary handling, day–count/calendar micro–conventions, smoothing/extrapolation of wings, and our *two–point* (variance/volatility) calibration which keeps our law of future realised variance *skewless* by design.

3.6 Variance–of–Variance (Vol-of-Var) vs. Benchmark

We now compare the *volatility of variance* parameter that drives the distribution of the future realised variance X in our model with the analogous quantity used in the benchmark pricer. In our framework this is the per-$\sqrt{\text{year}}$ log–volatility ν of X (cf. §2.3); the benchmark uses its own construction that embeds skew. For illustration we show one *uncapped* and one *capped* instrument.

Observations. Both paths co-move closely over time and react similarly to stress episodes. Our series is generally a touch higher—consistent with a *two–point* (variance/volatility) calibration that is *effectively skewless*, whereas the benchmark incorporates skew. Near shocks and close to maturity the benchmark occasionally adjusts more aggressively, but the resulting pricing impact remains small.

[1] Recall our calibration is pinned by the two quotes $(K_{\mathrm{var}}, K_{\mathrm{vol}})$ per maturity; this leaves no degree of freedom for distributional skew, hence our model remains *effectively skewless*.

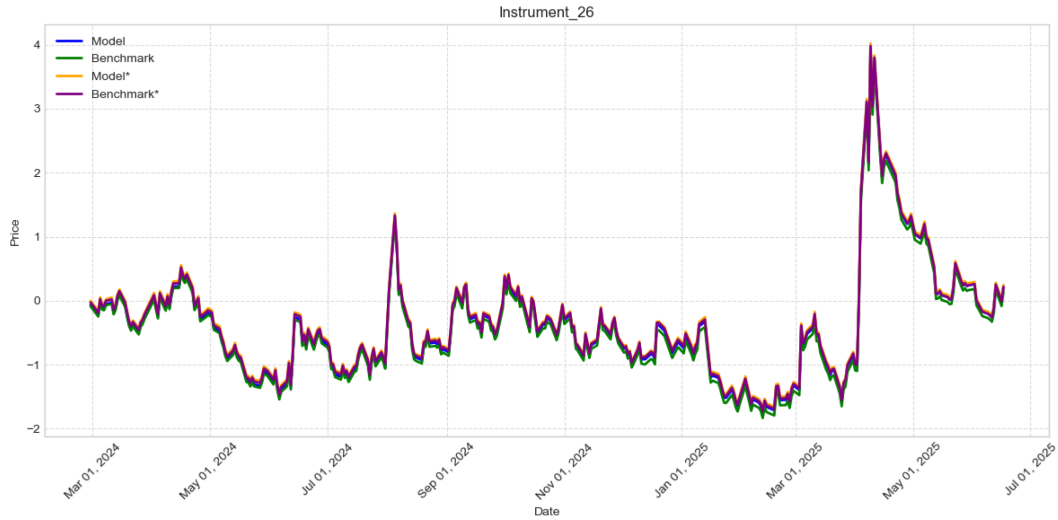

Figure 3.18: *
(a) Prices: model (blue), benchmark (green), model* (orange), benchmark* (purple).

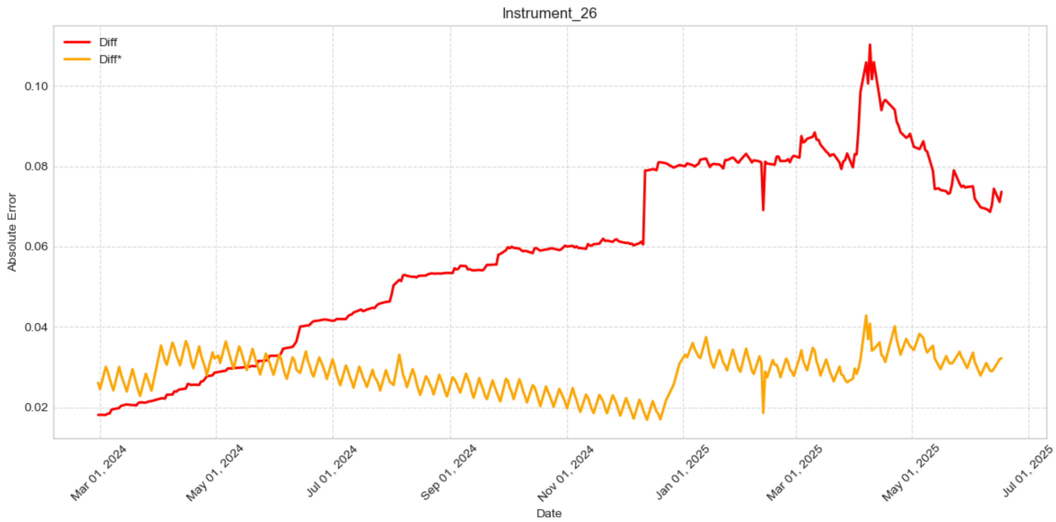

Figure 3.19: *
(b) Absolute errors: Diff (red) vs Diff* (orange).

Figure 3.20: Instrument 26 — **uncapped**: same–data / no–skew comparison.

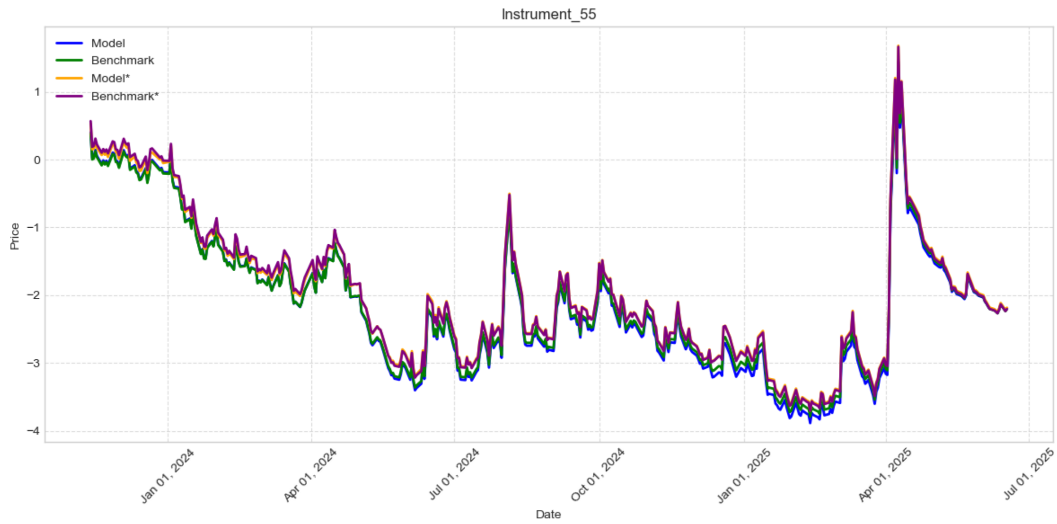

Figure 3.21: *
(a) Prices: model (blue), benchmark (green), model* (orange), benchmark* (purple).

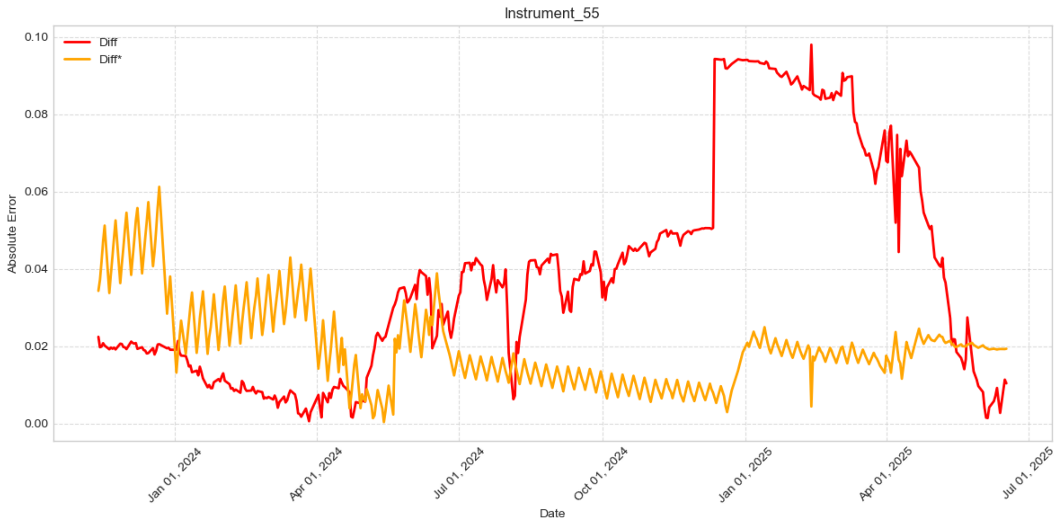

Figure 3.22: *
(b) Absolute errors: Diff (red) vs Diff* (orange).

Figure 3.23: Instrument 55 — **capped**: same–data / no–skew comparison.

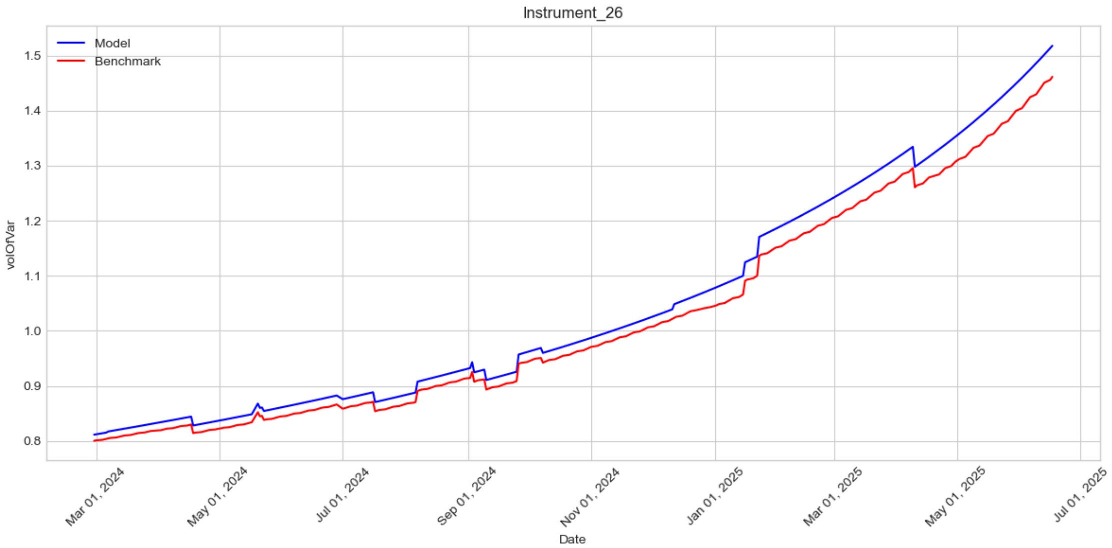

Figure 3.24: Instrument 26 — **uncapped** — time series of vol-of-var: model (blue) vs. benchmark (red).

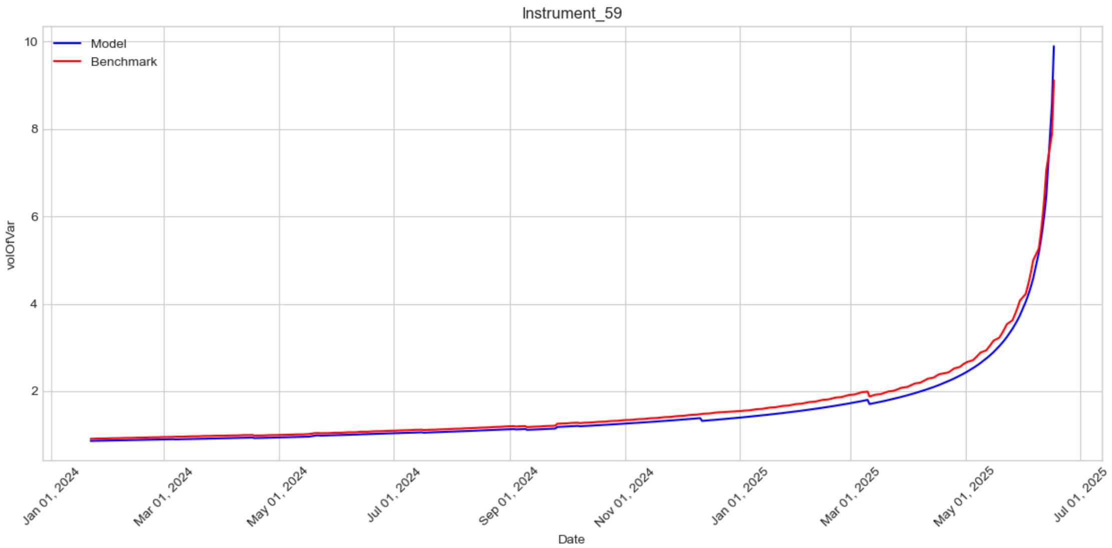

Figure 3.25: Instrument 59 — **capped** — time series of vol-of-var: model (blue) vs. benchmark (red).

Impact on pricing. The price discrepancies documented in §3.5 and §3.5 are *inferior to* 0.1 vol (ten bps in volatility units) across the sample. This magnitude is within desk tolerance for volatility–swap marks; traders are aware of this model/benchmark dispersion and routinely account for it in P&L explain and risk calculations. Importantly, these differences are *not* due to numerical convergence (we use a large grid N), but arise from market–data construction (capped/uncapped fair strikes) and the skew assumption (our skewless two–point calibration versus the benchmark's skewed dynamics).

Chapter 4

Conclusion and Outlooks

4.1 Summary of Contributions

This thesis develops a complete, desk–grade pipeline to value and hedge *seasoned, capped volatility swaps* using only vanilla options and CSA–consistent discounting. The main contributions are:

1. **Seasoned/capped functional replication on finite strikes.** We derive a Carr–Madan–type representation for the seasoned payoff $f(x) = \sqrt{ax+b}$ and show that a volatility cap induces an explicit boundary term (*CapAdj*); see Proposition 2.3.1 and Proposition 2.5.1. Continuous integrals are converted into a *finite–strike* implementation aligned with the law of future realised variance (Section 2.6).

2. **Parsimonious market model for future realised variance.** We model $X = \mathrm{RV}(t,T)$ as lognormal, calibrated *per maturity* from exactly two market points (variance– and volatility–swap fair strikes). Uncapped maturities admit a closed form for ν; capped maturities follow from a two–equation system in (μ,ν) (Section 2.3). The fitted law is *effectively skewless* but robust.

3. **Production–oriented numerics.** A split z–grid (puts left, calls right) targets probability mass and suppresses truncation error; for capped contracts the integral stops at the variance cutoff and the tail is handled exactly by CapAdj. Convergence tests (Section 3.2) show PV errors vanish as the grid is refined; production runs use large N.

4. **Collateral–consistent discounting.** All prices are computed in forward units and discounted under the trade's CSA currency using OIS curves (Section 3.3), keeping multi–currency books comparable and auditable.

5. **Empirical validation.** Across regimes, model PVs co–move tightly with a bank benchmark on capped and uncapped instruments (Section 3.5). Residual differences are traced to (i) construction of capped/uncapped fair strikes and (ii) the benchmark's skewed dynamics versus our two–point (skewless) law. A controlled *same–data / no–skew* experiment (Section 3.5) reduces the gap further. Remaining pricing differences are < 0.1 vol, operationally acceptable for trading and risk.

4.2 Practical Implications

- **Front–office usage.** Finite–strike baskets with explicit CapAdj translate directly into hedging tickets and XVA/risk explain. The seasoned structure $f(x) = \sqrt{ax+b}$ makes past/future and day–count effects explicit.

- **Risk & controls.** Deterministic weights, transparent truncation rules, and CSA–consistent discounting support daily P&L explain and model validation. Sensitivity to cap proximity is isolated in a single boundary term.

- **Run–time & scale.** The split–grid quadrature achieves high accuracy with modest strike counts; the replication layer is model–agnostic given an arbitrage–free vanilla surface.

4.3 Limitations

- **Skewlessness by design.** Pinning (μ, ν) from two quotes per maturity leaves no degree of freedom to fit skew/asymmetry of X. This is the main structural difference vs. a benchmark with skewed dynamics.

- **Surface and micro–conventions.** Results depend on construction of capped/uncapped fair strikes, wing smoothing, and calendars/day–count. Small convention mismatches can be amplified when caps are near binding.

- **Jump sensitivity.** Concavity ($\sqrt{\cdot}$) and caps increase sensitivity to jump episodes. Our static replication captures this via the option strip and CapAdj, but we do not model a jump component for X.

4.4 Outlooks and Future Work

(1) Enriching the law of future realised variance

- *Skew–aware X distributions.* Replace lognormal by shifted–lognormal, mixtures, or moment–matched families; or add one *skew* parameter per maturity. Calibrate with one extra market constraint (e.g., a wing–weighted moment from the vanilla strip or broker quotes for options on variance when available).

- *Jump overlays.* Superimpose a rare–jump component for X (compound–Poisson on $\ln X$) to better capture stress days; infer jump intensity from tail diagnostics on the surface.

- *Cross–maturity coupling.* Introduce a light prior or AR(1) regularisation on $\nu(T)$ to smooth the term structure over time while preserving transparency.

(2) Hedging, Greeks, and overlays

- *Static vs. dynamic mix.* Quantify the marginal benefit of a minimal dynamic overlay (delta/vega) on top of the static basket, especially near cap activation and around fixings.

- *Greeks to surface shocks.* Differentiate the finite–strike functional w.r.t. local implied vol nodes to obtain explainable bucketed sensitivities; compare with the benchmark's risk attributions.

- *Jump–robust hedges.* Explore targeted tail hedges (deep OTM puts/calls) that reduce model–form error during volatility jumps while leaving carry largely unaffected.

(3) Stress testing and robustness

- *Historical scenarios.* Replay 2018 Volmageddon, COVID–19 (2020), and 2025 tariff headlines using the same replication to benchmark realised performance of capped vs. uncapped structures.
- *Wing sparsity.* Systematically test truncation and extrapolation schemes under sparse wings; document error bars vs. bid–ask.

(4) Beyond equity underlyings

- *Single–names and other asset classes.* Apply the same finite–strike framework to single–name equities and to FX/rates where definitions and fixing rules differ; catalogue any asset–class–specific adjustments.

4.5 Closing Remarks

The proposed framework connects clean theory (Carr–Madan functional replication) to a practical, auditable implementation for seasoned and capped volatility swaps. Its strengths are transparency, stability, and direct tradability of the hedges. Residual spreads to a sophisticated benchmark are small (< 0.1 vol) and well explained by market–data construction and skew assumptions. The suggested extensions—chiefly a minimally parametric skew for X, jump overlays, and improved fair–strike harmonisation—offer a clear path to further reduce these differences while preserving the simplicity that makes the engine robust in production.

Appendix A

Technical Proofs

A.1 Proof of the Black formula

Proof. Assume $\ln Y \sim \mathcal{N}(m, v)$ with $v = \sigma^2\tau$ and choose $m = \ln F - \frac{1}{2}v$ so that $\mathbb{E}[Y] = e^{m+\frac{1}{2}v} = F$. Write $Y = e^X$ with $X = m + \sqrt{v}\,Z$, $Z \sim \mathcal{N}(0,1)$. Then

$$\mathbb{E}\big[(Y-K)^+\big] = \mathbb{E}\big[e^X \mathbf{1}_{\{X > \ln K\}}\big] - K\,\mathbb{P}(X > \ln K).$$

For the second term,

$$\mathbb{P}(X > \ln K) = \Phi\left(\frac{m - \ln K}{\sqrt{v}}\right) = \Phi\left(\frac{\ln(F/K) - \frac{1}{2}v}{\sqrt{v}}\right) = \Phi(d_0),$$

since $d_0 := \dfrac{\ln(F/K) - \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}}$.

For the first term, use the standard Gaussian tilting identity $e^{az-\frac{1}{2}a^2}\varphi(z) = \varphi(z-a)$ (with $a = \sqrt{v}$), where φ is the standard normal pdf. With the change of variable $z = (x-m)/\sqrt{v}$ we get

$$\begin{aligned}\mathbb{E}\big[e^X \mathbf{1}_{\{X > \ln K\}}\big] &= \int_{\ln K}^{\infty} e^x \frac{1}{\sqrt{v}}\,\varphi\left(\frac{x-m}{\sqrt{v}}\right)\,dx = e^{m+\frac{1}{2}v}\int_{\frac{\ln K - m}{\sqrt{v}}}^{\infty} \varphi(z - \sqrt{v})\,dz \\ &= e^{m+\frac{1}{2}v}\,\Phi\left(\sqrt{v} - \frac{\ln K - m}{\sqrt{v}}\right).\end{aligned}$$

Since $e^{m+\frac{1}{2}v} = F$ and $v = \sigma^2\tau$, this equals

$$F\,\Phi\left(\frac{\ln(F/K) + \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}}\right) = F\,\Phi(d_1),$$

with $d_1 := \dfrac{\ln(F/K) + \frac{1}{2}\sigma^2\tau}{\sigma\sqrt{\tau}}$.

Combining both terms,

$$\mathbb{E}\big[(Y-K)^+\big] = F\,\Phi(d_1) - K\,\Phi(d_0),$$

and, by construction, $d_0 = d_1 - \sigma\sqrt{\tau}$. This is the Black formula. $\square$

A.2 Proof of Theorem 2.4.1 (Carr–Madan representation)

Proof. Let $F = \mathbb{E}[X]$ and define, for $x \geq 0$,

$$H(x) \;:=\; f(F) + f'(F)(x - F) + \int_0^F f''(k)\,(k-x)^+\,dk + \int_F^\infty f''(k)\,(x-k)^+\,dk.$$

By the assumptions on f and f', and the stated integrability of f'' against put/call premia, both integrals are finite for each $x \geq 0$, so H is well-defined.

Differentiate under the integral sign (Leibniz/Dominated Convergence; note $\partial_x(k-x)^+ = -\mathbf{1}_{\{x<k\}}$ and $\partial_x(x-k)^+ = \mathbf{1}_{\{x>k\}}$):

$$H'(x) = f'(F) - \int_0^F f''(k)\,\mathbf{1}_{\{x<k\}}\,dk + \int_F^\infty f''(k)\,\mathbf{1}_{\{k<x\}}\,dk = f'(F) + \int_F^x f''(k)\,dk.$$

Hence $H''(x) = f''(x)$ for all $x > 0$, and $H(F) = f(F)$, $H'(F) = f'(F)$. Thus $H \equiv f$ on $(0, \infty)$, yielding the pointwise identity

$$f(x) = f(F) + f'(F)(x - F) + \int_0^F f''(k)\,(k-x)^+\,dk + \int_F^\infty f''(k)\,(x-k)^+\,dk. \quad \text{(A.2.1)}$$

Taking expectations in (A.2.1) with $x = X$ and using $\mathbb{E}[X - F] = 0$, and Tonelli/Fubini (justified by the finiteness of $\int_0^F |f''(k)|\,\mathrm{Put}(k)\,dk$ and $\int_F^\infty |f''(k)|\,\mathrm{Call}(k)\,dk$), we obtain

$$\mathbb{E}[f(X)] = f(F) + \int_0^F f''(k)\,\mathrm{Put}(k)\,dk + \int_F^\infty f''(k)\,\mathrm{Call}(k)\,dk,$$

which is (2.4.1). □

Remark. Centering at any $K > 0$ gives $\mathbb{E}[f(X)] = f(K) + f'(K)(F-K) + \int_0^K f''(k)\,\mathrm{Put}(k)\,dk + \int_K^\infty f''(k)\,\mathrm{Call}(k)\,dk$; taking $K = F$ eliminates the linear term.

Appendix B

Full Instruments Dataset

Convention. Throughout the dataset, we encode *uncapped* instruments by setting `Cap`=0, i.e., no truncation is applied to the realised volatility/variance.

Table B.1: Instruments dataset

Instrument	Strike	t0	Maturity	Cap
1	0.197804	10/06/2025	15/12/2028	0
2	0.1875	20/05/2025	18/12/2026	0
3	0.1825	04/06/2025	18/12/2026	0.45625
4	0.191998	27/12/2024	18/12/2026	0
5	0.1735	11/02/2025	18/12/2026	0.43375
6	0.194	31/03/2025	18/12/2026	0
7	0.1703	29/01/2025	19/06/2026	0.4258
8	0.1714	06/02/2025	19/06/2026	0.4285
9	0.1827	05/03/2025	19/06/2026	0.4568
10	0.173	26/03/2025	19/06/2026	0.4325
11	0.1845	10/03/2025	19/06/2026	0.4612
12	0.188	13/05/2025	19/06/2026	0.47
13	0.1839	28/05/2025	19/06/2026	0.4598
14	0.19	31/03/2025	20/03/2026	0
15	0.1945	12/03/2025	20/03/2026	0
16	0.228	04/04/2025	20/03/2026	0
17	0.2195	24/04/2025	20/03/2026	0
18	0.1674	03/09/2024	19/12/2025	0.4185
19	0.1709	26/06/2024	19/12/2025	0.42725
20	0.185	13/01/2025	19/12/2025	0.4625
21	0.1795	15/01/2024	19/12/2025	0.4488
22	0.1897	03/01/2024	19/12/2025	0.4743
23	0.1722	31/01/2024	19/12/2025	0.4305
24	0.1721	21/02/2024	19/12/2025	0.43
25	0.1667	22/02/2024	19/12/2025	0.4168
26	0.1765	29/02/2024	19/12/2025	0
27	0.17	06/03/2024	19/12/2025	0.425
28	0.1701	03/04/2024	19/12/2025	0.4253
29	0.1689	08/05/2024	19/12/2025	0.4223
30	0.1663	03/06/2024	19/12/2025	0.4158
31	0.1666	04/06/2024	19/12/2025	0.4165
32	0.165	13/06/2024	19/12/2025	0.4125
33	0.1648	04/06/2024	19/12/2025	0.412
34	0.1702	20/06/2024	19/12/2025	0.4255

(continues on next page)

(continued from previous page)				
Instrument	Strike	t0	Maturity	Cap
35	0.179	02/08/2024	19/12/2025	0.4475
36	0.1679	15/07/2024	19/12/2025	0.41975
37	0.1679	17/07/2024	19/12/2025	0.4198
38	0.166	16/08/2024	19/12/2025	0.415
39	0.183	07/08/2024	19/12/2025	0.4575
40	0.169	26/08/2024	19/12/2025	0.4225
41	0.1795	05/11/2024	19/12/2025	0.44875
42	0.1711	14/11/2024	19/12/2025	0.4278
43	0.1723	12/11/2024	19/12/2025	0.4308
44	0.174566	14/11/2024	19/12/2025	0
45	0.184	13/01/2025	19/12/2025	0.46
46	0.164	22/01/2025	19/12/2025	0.41
47	0.1618	22/01/2025	19/12/2025	0.4045
48	0.1705	27/01/2025	19/12/2025	0.42625
49	0.1664	11/02/2025	19/12/2025	0.416
50	0.1674	28/02/2025	19/12/2025	0.4185
51	0.194	31/03/2025	19/12/2025	0
52	0.1616	23/01/2025	19/12/2025	0.404
53	0.1674	03/09/2024	19/12/2025	0.4185
54	0.172	24/10/2024	19/09/2025	0.43
55	0.1769	13/11/2023	20/06/2025	0.4423
56	0.1841	08/11/2023	20/06/2025	0.4603
57	0.1809	04/01/2024	20/06/2025	0.452
58	0.1809	04/01/2024	20/06/2025	0.452
59	0.1659	22/01/2024	20/06/2025	0.4148
60	0.1658	20/02/2024	20/06/2025	0.4145
61	0.167	14/03/2024	20/06/2025	0.4175
62	0.166	15/03/2024	20/06/2025	0.415
63	0.1657	20/03/2024	20/06/2025	0.41425
64	0.173	27/06/2024	20/06/2025	0.4325
65	0.177	11/09/2024	20/06/2025	0.4425
66	0.162	21/01/2025	20/06/2025	0

Bibliography

Breeden, D. T. and R. H. Litzenberger (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business* **51**(4), 621–651.

Carr, P. and D. B. Madan (1998). Towards a theory of volatility trading. In R. Jarrow (Ed.), *Volatility: New Estimation Techniques for Pricing Derivatives*, 417–427. London: Risk Publications.

Carr, P. and D. B. Madan (1999). Option valuation using the fast Fourier transform. *Journal of Computational Finance* **2**(4), 61–73.

Carr, P. and R. Lee (2008). Volatility derivatives. *Annual Review of Financial Economics* **1**, 319–339.

Carr, P. and L. Wu (2006). A tale of two indices. *Journal of Derivatives* **13**(3), 13–29.

Demeterfi, K., E. Derman, M. Kamal, and J. Zou (1999). More than you ever wanted to know about volatility swaps. *Goldman Sachs Quantitative Strategies Research Notes.*

Broadie, M. and A. Jain (2008). Pricing and hedging volatility derivatives. *Journal of Derivatives* **15**(3), 7–24.

Bossu, S., E. Strasser, and R. Guichard (2005). Just what you need to know about variance swaps. JPMorgan Equity-Derivatives Research Note.

Bossu, S. (2019). *Variance Swaps Revisited: New Insights on Hedging and Pricing.* Working paper, Université Paris-Dauphine.

Neuberger, A. (1994). The log contract. Working paper, London Business School.

Dupire, B. (1994). Pricing with a smile. *Risk* **7**(1), 18–20.

Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide.* Wiley Finance.

Gatheral, J. and A. Jacquier (2014). Arbitrage-free SVI volatility surfaces. *Quantitative Finance* **14**(1), 59–71.

Lee, R. W. (2004). The moment formula for implied volatility at extreme strikes. *Mathematical Finance* **14**(3), 469–480.

Fengler, M. (2009). Arbitrage-free smoothing of the implied volatility surface. *Quantitative Finance* **9**(4), 417–428.

Cont, R. and P. Tankov (2004). *Financial Modelling with Jump Processes.* Chapman & Hall/CRC.

Piterbarg, V. (2010). Funding beyond discounting: collateral agreements and derivatives pricing. *Risk* **23**(2), 97–102.

Henrard, M. (2014). *Interest Rate Modelling in the Multi-Curve Framework: Foundations, Evolution, and Implementation.* Palgrave Macmillan.

International Swaps and Derivatives Association (2011). *2009/2011 ISDA Equity Variance Swap Definitions.* ISDA Publications.

Bühler, H. (2022). Empirical observations on variance-of-variance surfaces in equity markets. SSRN Working Paper #4071234.

Lee, P. (2025). Optimal variance swap replication. SSRN Working Paper #5047932.