Pricing Constant Maturity Swaps with Gaussian Short-Rate and Quasi-Gaussian Stochastic-Volatility Models

Matias Koponen

School of Science

Thesis submitted for examination for the degree of Master of Science in Technology.
Helsinki 12.05.2026

Supervisor

Professor Ahti Salo

Advisor

Veino Ville, Ph.D.

Aalto University
School of Science

Aalto University
School of Science

Aalto University, P.O. BOX 11000, 00076 AALTO
www.aalto.fi
Abstract of the master's thesis

Author Matias Koponen

Title Pricing Constant Maturity Swaps with Gaussian Short-Rate and Quasi-Gaussian Stochastic-Volatility Models

Degree programme Master's Program in Mathematics and Operations Research

Major Applied Mathematics

Supervisor Professor Ahti Salo

Advisor Veino Ville, Ph.D.

Date 12.05.2026 **Number of pages** 85+17 **Language** English

Abstract

This thesis studies the pricing of Constant Maturity Swap (CMS) instruments in a post-crisis multi-curve interest rate framework. The focus is on callable CMS products, whose valuation combines CMS convexity adjustments, swaption smile effects, spread correlation, and Bermudan exercise decisions.

The thesis develops a Monte Carlo workflow for calibrating and pricing a multi-factor quasi-Gaussian model with stochastic volatility (QG-SV), using the Andersen–Piterbarg benchmark-rate framework. The model is calibrated to EUR swaption normal implied volatility data and compared with Gaussian and skew-only model restrictions.

In the empirical study, the two-factor stochastic-volatility model reproduces the 17 March 2026 EUR swaption smile with a root-mean-square error of 2.00 basis points in normal implied volatility. The results show that non-callable CMS valuations are largely determined by analytical CMS replication inputs, whereas callable CMS products are materially affected by the model used to simulate future interest-rate dynamics and exercise decisions. In a stress-amplified callable CMS spread swap, displacement skew increases the callable value relative to the Gaussian baseline, while stochastic volatility partly offsets this increase.

Overall, the full quasi-Gaussian stochastic-volatility specification gives the most accurate EUR swaption smile fit among the tested models. Callable CMS valuation is driven by the interaction of smile fit, factor structure, and Bermudan exercise policy.

Keywords CMS instruments, multi-curve framework, quasi-Gaussian models, stochastic volatility, Monte Carlo calibration, callable CMS spread, convexity adjustment, benchmark rate parametrization

Aalto University
School of Science

Aalto-yliopisto, PL 11000, 00076 AALTO
www.aalto.fi
Diplomityön tiivistelmä

Tekijä Matias Koponen

Työn nimi Constant Maturity Swap -instrumenttien hinnoittelu: gaussisten lyhyen koron mallien ja stokastisen volatiliteetin kvasi-gaussisten mallien vertailu

Koulutusohjelma Master's Program in Mathematics and Operations Research

Pääaine Applied Mathematics

Työn valvoja Professor Ahti Salo

Työn ohjaaja Veino Ville, Ph.D.

Päivämäärä 12.05.2026 **Sivumäärä** 85+17 **Kieli** Englanti

Tiivistelmä

Diplomityössä tarkastellaan Constant Maturity Swap (CMS) -instrumenttien hinnoittelua finanssikriisin jälkeisessä monikäyräisessä korkokehikossa. Painopiste on irtisanottavissa CMS-tuotteissa, joissa yhdistyvät CMS-konveksisuus, swaptioiden volatiliteettihymy, spread-korrelaatio ja Bermudan-tyyppinen irtisanominen.

Työssä kehitetään Monte Carlo -menetelmä monifaktorisen kvasi-gaussisen stokastisen volatiliteetin mallin kalibrointiin ja hinnoitteluun Andersenin ja Piterbargin viitekorkokehikossa. Malli kalibroidaan EUR-swaptioiden normaalivolatiliteetteihin ja sitä verrataan gaussiseen ja vinousrajoitteiseen malliin.

Empiirisessä tutkimuksessa kaksifaktorinen stokastisen volatiliteetin malli sovittaa 17.3.2026 EUR-swaptioiden volatiliteettihymyn 2,00 korkopisteen keskineliövirheen neliöjuurella. Ei-irtisanottavat CMS-arvot määräytyvät pitkälti analyyttisen CMS-replikaation syötteistä, kun taas irtisanottavat CMS-tuotteet riippuvat korkodynamiikkaa ja irtisanomispäätöksiä simuloivasta mallista. Stressatussa CMS-spread-swapissa siirtymävinous kasvattaa arvoa, ja stokastinen volatiliteetti vaimentaa osan kasvusta.

Kokonaisuutena täysi kvasi-gaussinen stokastisen volatiliteetin malli sovittaa EUR-swaptioiden volatiliteettihymyn tarkimmin. Irtisanottavien CMS-tuotteiden arvostus riippuu volatiliteettihymyn, faktorirakenteen ja Bermudan-tyyppisen irtisanomisstrategian yhteisvaikutuksesta.

Avainsanat CMS-instrumentit, monikäyräinen korkokehikko, kvasi-gaussiset mallit, stokastinen volatiliteetti, Monte Carlo -kalibrointi, irtisanottavat CMS-tuotteet

Acknowledgements

I would like to express my gratitude to those who supported me during the preparation of this thesis. I am especially thankful to my supervisor Ahti Salo for supervising this thesis and to my advisor Ville Veino for his expertise and patience during the writing process. Moreover, I wish to thank my colleague Vesa Tähtinen for his substantial guidance on the theoretical framework, Jaakko Juntunen for providing the initial topic, and many other colleagues along the way. Their encouragement made it possible to complete this thesis alongside the many practical and technical challenges involved in the project. I am also thankful to my friends and family for their support, patience, and perspective during the writing process. Lastly, I wish to acknowledge and extend my gratitude to my wife Maria, whose support, patience and love were invaluable throughout this journey.

Helsinki, May 12 2026

Matias Koponen

Contents

Notation

Table 0.1: Principal notation used throughout the thesis.

Symbol	Description
$P(t,T)$	OIS discount bond price
$f(t,T)$	Instantaneous forward rate
$r(t)$	Short rate, $r(t) = f(t,t)$
d	Number of rate factors
$S_{n,m}(t)$	Forward swap rate (start T_n, maturity m years)
$A(t)$	Annuity (swap-rate numeraire)
τ_i	Fixed-leg accrual fraction
δ_i	Floating-leg accrual fraction
$\mathbb{Q}$	Risk-neutral measure
$\mathbb{Q}^T$	T-forward measure
$\mathbb{Q}^A$	Annuity measure
Quasi-Gaussian model	
$x_k(t)$	Rate factor k (state variable)
$y_{kl}(t)$	Cross-variance of factors k, l
$z(t)$	CIR variance process (stochastic volatility)
κ_k	Mean reversion speed of factor k
$G_k(t,T)$	Accrual factor, $\int_t^T e^{-\int_t^u \kappa_k\, ds}\, du$
$h(t)$	Maturity function, $h(t) = e^{-\kappa t}$ (constant κ)
$\lambda_k(t)$	Base volatility (scaling), also denoted $\sigma^0(t)$
$b_k(t)$	Displacement parameter (skew control)
η	QG-SV vol-of-vol (smile curvature), also denoted ν; η_G denotes the second G2++ factor volatility
θ	CIR mean reversion speed for $z(t)$
$\bar{z}$	Long-run variance level
$\sigma_x(t)$	State variable diffusion matrix
$\sigma_k^f(t)$	Benchmark forward rate volatility
H^f	Benchmark shift matrix, $H^f_{ik} = e^{-\kappa_k \Delta_i}$
D^f	Cholesky factor of benchmark correlation
ρ_{kl}	Inter-factor correlation
Δ_k	Benchmark tenor assigned to factor k
Calibration	
RMS	Root mean square error (in bp)
ATM	At-the-money strike
SA	Semi-analytical (frozen-sensitivity approximation)
MC	Monte Carlo simulation

Standing Assumptions

The curve and collateral terms used in these assumptions are defined formally in Chapter 2; they are stated here first so that all subsequent model and pricing results share the same probability-space and admissibility conditions.

Assumption 0.1 (Market model and admissibility)**.** All stochastic processes are defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{Q})$ satisfying the usual conditions. The measure $\mathbb{Q}$ is the collateralised risk-neutral measure associated with the OIS money-market account. Discount factors $P^d(t,T)$ are strictly positive and continuously differentiable in maturity, and all numeraires used for measure changes are strictly positive collateralised value processes. For each such numeraire N, the discounted process N_t/B_t^d is assumed to be a true $\mathbb{Q}$-martingale on the finite horizon considered, so that the associated numeraire measure is equivalent to $\mathbb{Q}$.

Assumption 0.2 (Curves and collateralisation)**.** The EUR market is represented by an OIS discount curve and tenor-specific forwarding curves. Basis spreads between forwarding and discounting curves are deterministic functions of time in the numerical implementation. This is the standard multi-curve approximation used in the empirical study and isolates the model comparison to interest-rate volatility, correlation, and exercise dynamics.

Assumption 0.3 (Regularity of model coefficients)**.** The volatility and displacement functions appearing in the Gaussian, QG, and QG-SV models are Borel measurable in time, locally Lipschitz in the state variables, and satisfy linear-growth bounds on compact time intervals. The CIR variance process is interpreted with the usual non-negative square-root diffusion; when the Feller condition is not imposed numerically, the moment-matched CIR step of Chapter 6 defines the Monte Carlo approximation.

Assumption 0.4 (Integrability and martingale conditions)**.** All discounted payoff processes used in the pricing identities are integrable, and, for every numeraire N used below, the ratios X_T/N_T appearing in conditional expectations are integrable under the corresponding numeraire measure. The density processes used for measure changes and the stochastic exponentials generated by the QG/QG-SV bond dynamics are assumed to be true martingales rather than merely local martingales. The interchanges of integration and expectation in the replication formulae are justified by the same payoff-integrability conditions, for instance by absolute integrability of the Carr–Madan replication integrals. Continuous-time pricing identities are therefore separated from the Monte Carlo discretisation used to approximate them numerically.

1 Introduction

1.1 Background and Motivation

The global financial crisis of 2007–2009 changed the way interest rate derivatives are priced. Significant basis spreads appeared between formerly fungible tenor rates, most notably the divergence between EURIBOR rates of different tenors and the OIS discounting rate, making the pre-crisis single-curve framework unsuitable for collateralised multi-tenor valuation. In its place, the *multi-curve framework* became the new paradigm, requiring separate estimation of discount curves and forward rate curves for each tenor [1, 7, 33]. This transition, while conceptually straightforward, added considerable complexity to the calibration and pricing of interest rate derivatives.

Among the products most affected by this shift are *Constant Maturity Swap* (CMS) instruments. CMS swaps, CMS spread swaps, and their callable variants form an important part of the EUR interest rate derivatives market. A CMS coupon pays the prevailing swap rate of a specified tenor on each fixing date, providing exposure to the level and shape of the yield curve. CMS spread coupons, which pay the difference between two swap rates of different tenors, offer direct exposure to the slope of the term structure. Callable versions of these products (Bermudan-style instruments with issuer exercise rights) add a further layer of complexity through the embedded optionality.

Pricing CMS-linked instruments is difficult for several reasons. First, the CMS coupon itself requires a *convexity adjustment*: the expected value of a future swap rate under the payment measure differs from its forward value due to the nonlinear relationship between swap rates and discount factors [29]. Second, CMS spread options depend on the *correlation* between the underlying swap rates, which is governed by the factor structure of the chosen interest rate model. Third, callable CMS products require the valuation of Bermudan exercise rights, demanding either regression-based Monte Carlo techniques [46] or semi-analytical approximations.

Traditional Gaussian short-rate models (the Hull–White model [39] and its multi-factor extension G2++) have long served as workhorses for callable product pricing. They offer closed-form bond pricing, efficient calibration to the at-the-money swaption surface, and tractable Bermudan exercise evaluation. However, Gaussian dynamics imply a flat volatility smile, whereas EUR swaption markets display skew and curvature that are commonly modelled through displaced-diffusion, SABR, local-volatility or stochastic-volatility mechanisms [30, 58, 5]. For CMS instruments,

which are sensitive to the distribution of future swap rates beyond the at-the-money volatility, ignoring this smile information can materially affect valuation.

The *quasi-Gaussian* (QG) framework, developed by [5], addresses this limitation through two extensions. First, a *local volatility* specification, typically of displaced-diffusion or CEV type, introduces skew into the swap rate distribution. Second, a *stochastic volatility* component, modelled as a CIR process driving the instantaneous variance of the short rate, generates the smile curvature observed in market prices. The resulting QG-SV model retains the Markovian state-space structure needed for efficient simulation while being flexible enough to match the observed swaption smile.

In this thesis we develop a pricing framework for CMS instruments, from vanilla CMS swaps through callable CMS spread structures, within the multi-factor quasi-Gaussian stochastic volatility model. The framework is realised as a reproducible Monte Carlo methodology, providing a transparent basis for the empirical analysis.

1.2 Research Questions

This thesis addresses three interrelated research questions:

1. Under a common multi-curve numerical setup, how strongly do factor structure, skew, and stochastic volatility change the pricing of European swaptions and callable CMS products, especially CMS spread structures?

2. Can the multi-factor QG-SV model be calibrated with the same Monte Carlo machinery used for downstream callable pricing, so that vanilla fit and exotic valuation remain numerically consistent?

3. Once a multi-factor smile model is calibrated, how materially do higher-order rate factors beyond the dominant level and slope modes affect callable pricing?

1.3 Contributions

This thesis makes the following contributions:

1. We develop a reproducible numerical methodology for the multi-factor QG-SV model, including Monte Carlo pricing methods for European and Bermudan CMS-linked products and a Monte Carlo bootstrap calibration workflow. The framework supports an arbitrary number of factors and identifies conditions under which semi-analytical calibration based on frozen sensitivities can be inconsistent with full Monte Carlo pricing.

2. We design a controlled empirical comparison protocol for the thesis results chapter, where Table A benchmarks LGM, G2++, and QG-SV on an ATM Bermudan anchor; Table B compares nested Gaussian, QG+Skew, and QG-SV variants under the same full-smile calibration family on the full callable CMS portfolio; and Table C interprets factor sensitivity through the PCA of the calibrated multi-factor correlation structure.

3. We obtain three case-study findings using EUR market data as of 17 March 2026 and relative-tenor trades rebased to that date: multi-speed Gaussian dynamics materially affect Bermudan exercise premia, improved full-smile fit does not translate monotonically into every callable price, and low-variance residual factor modes can still move Bermudan valuations materially once calibration and exercise interact.

1.4 Thesis Structure

The thesis is organised as follows. Chapter 2 reviews the post-crisis multi-curve framework, establishes the distinction between discount and forwarding curves, and presents the bootstrapping methodology for OIS, EURIBOR swaps, and basis swaps. Chapter 3 defines the CMS instruments analysed in this thesis: CMS swaps, CMS caps and floors, CMS spread options, and callable CMS spread swaps, together with their market conventions and payoff structures.

Chapter 4 develops the interest rate models used in the analysis, progressing from the one-factor Gaussian short-rate model (Hull–White/GSR), through the multi-factor Gaussian model (G2++), to the quasi-Gaussian framework with local volatility (QG-LV) and stochastic volatility (QG-SV). For each model the chapter presents the dynamics, derives the bond pricing formulas, and discusses calibration. Chapter 5 then presents the pricing techniques for CMS instruments, including the convexity adjustment for CMS coupons, the replication approach for CMS caplets and floorlets, and the regression-based Monte Carlo method for Bermudan exercise evaluation.

Chapter 6 describes the numerical methodology: the Monte Carlo simulation scheme, the bootstrap calibration algorithm, the Levenberg–Marquardt optimisation procedure [44, 49], and the parameter-restriction protocol that enables controlled model comparison. Chapter 7 is the empirical core of the thesis. It is organised around three increasingly targeted comparisons, namely Table A, an ATM cross-family benchmark across LGM, G2++, and QG-SV; Table B, a nested full-smile comparison

inside the Gaussian/QG+Skew/QG-SV family on the full callable CMS portfolio; and Table C, a PCA-guided factor study for the multi-factor QG-SV specification. Supporting sections then discuss the LGM benchmark caveat for spread-contingent callables, the semi-analytical versus Monte Carlo gap, and computational cost.

Chapter 8 summarises the findings, discusses practical implications, and outlines directions for future research. Appendix A contains technical proofs and derivations, and Appendix B provides supplementary calibration results and parameter tables.

1.5 Related Work

The theoretical foundation for this thesis rests on the finite-dimensional HJM and Cheyette-model lineage. Cheyette's Markovian term-structure construction [17], subsequent finite-dimensional HJM restrictions [64, 15], and the geometric finite-dimensional-realisation theory of Björk–Christensen and Björk–Svensson [8, 9] provide the mathematical background for the quasi-Gaussian representation. The treatment of quasi-Gaussian models by Andersen and Piterbarg [5] provides the notation and model specification adopted here. The multi-factor Markovian representation and the associated semi-analytical swaption pricing formulas are central to the numerical study.

The Gaussian short-rate benchmarks used for comparison descend from the affine short-rate tradition initiated by Vasicek and Ho–Lee and extended by Hull–White [70, 36, 37]. The foundational no-arbitrage curve framework builds on the Heath–Jarrow–Morton approach [32]. Market-model alternatives for directly modelling LIBOR and swap rates include the Brace–Gatarek–Musiela, Miltersen–Sandmann–Sondermann and Jamshidian frameworks [10, 52, 40]. The multi-curve framework, which became standard after the financial crisis, is developed in [33] and [27]. The practical aspects of curve construction, including the treatment of basis spreads and OIS discounting, draw on [1] and [21].

For CMS pricing specifically, the seminal convexity adjustment approach of Hagan [29] remains influential, while the replication and measure-change literature of Hunter–Jäckel, Pelsser, Pugachevsky and Cedervall–Piterbarg [38, 56, 61, 16] gives the foundation for the one-dimensional CMS formulas used as benchmarks. Mercurio [50] extends CMS and swap-rate formulae to the multi-curve setting, and Andersen–Piterbarg [4] gives the replication-based treatment followed in Chapter 5.

The gap addressed by this thesis is narrower than the general CMS literature. Static replication and SABR-style smile adjustments are well suited to plain CMS coupons and caplets, but callable CMS spread products require a joint dynamic

model for multiple swap rates at future exercise dates. The present work therefore studies a reproducible calibration and pricing workflow for a dynamic QG-SV model, and compares how factor structure, skew, and stochastic volatility affect callable CMS valuations in that setting.

The numerical methods employed in this thesis draw on established techniques: the Longstaff–Schwartz algorithm for American-style option pricing via regression-based Monte Carlo [46], the dual and error-analysis literature for simulation-based exercise methods [65, 3, 25], the Levenberg–Marquardt algorithm for nonlinear least-squares calibration [54], and Sobol quasi-Monte Carlo simulation for multi-factor models [67, 24]. Multilevel Monte Carlo [23] is discussed only as a possible acceleration technique in Chapter 7.

Numerical reproducibility is addressed by reporting the model parametrisation, calibration objective, simulation grid, path counts, optimisation settings, and calibration diagnostics. The thesis presentation is therefore kept at the level of models, algorithms, calibration objectives, and empirical results.

2 The Multi-Curve Framework

2.1 Introduction

The purpose of this chapter is to develop the theoretical foundation for pricing CMS instruments in the post-crisis multi-curve environment. The financial crisis of 2007–2009 made the assumption of a single universal discount curve untenable for collateralised multi-tenor valuation. Basis spreads between different LIBOR tenors and between LIBOR and OIS rates widened dramatically; the 3-month LIBOR–OIS spread reached historically extreme levels in September 2008, while tenor basis spreads also became economically significant [12, 59, 27].

The modern approach in interest rate modeling is to separate curves used for discounting and forward rate projection. The approach has been developed by Piterbarg [59], Bianchetti [7], and Henrard [33]. We will present the framework focusing on its application to CMS pricing, where convexity adjustments turn out to be highly sensitive to the multi-curve structure.

2.2 Pre-Crisis Single Curve Framework

In the pre-crisis world, practitioners assumed a single discount curve $P(t,T)$ for both discounting and forward rate projection. Here t is the valuation time, T a future date, $P(t,T)$ the time-t discount factor to T, and δ the accrual year fraction. LIBOR denotes the historical London Interbank Offered Rate; OIS rates are overnight-indexed swap rates based on risk-free rates (RFRs). The simply compounded LIBOR forward rate was computed as

$$L(t,T,T+\delta) = \frac{1}{\delta}\left(\frac{P(t,T)}{P(t,T+\delta)} - 1\right) \tag{2.1}$$

The par swap rate making a fixed-for-floating swap have zero value at inception was

$$S(t,t_0,t_n) = \frac{P(t,t_0) - P(t,t_n)}{\sum_{i=1}^{n} \tau_i P(t,t_i)} \tag{2.2}$$

This framework rested on the assumption that LIBOR represented a risk-free rate and that different LIBOR tenors were interchangeable. The crisis made these assumptions unsuitable for post-crisis collateralised valuation.

2.3 Multi-Curve Framework

2.3.1 Separation of Discounting and Forwarding

The fundamental principle is to use a risk-free discount curve $P^{\text{OIS}}(t,T)$ for discounting, while maintaining tenor-specific index curves $P^k(t,T)$ for forward rate projection. Here k labels the tenor and δ^k is its accrual fraction

$$L^k(t,T,T+\delta^k) = \frac{1}{\delta^k}\left(\frac{P^k(t,T)}{P^k(t,T+\delta^k)} - 1\right) \tag{2.3}$$

The OIS discount curve is constructed from overnight indexed swap contracts. For collateralized derivatives, the discount curve reflects the collateral remuneration rate.

2.3.2 Swap Pricing and Par Rates

A fixed-for-floating swap paying 3-month LIBOR, with coupon accrual fractions τ_i and fixed coupon rate c, has value

$$V_{\text{swap}}(t) = \sum_{i=1}^{n} \tau_i L^{3M}(t,t_{i-1},t_i) P^{\text{OIS}}(t,t_i) - c\sum_{i=1}^{n} \tau_i P^{\text{OIS}}(t,t_i) \tag{2.4}$$

The par swap rate is

$$S^{3M}(t,t_0,t_n) = \frac{\sum_{i=1}^{n} \tau_i L^{3M}(t,t_{i-1},t_i) P^{\text{OIS}}(t,t_i)}{\sum_{i=1}^{n} \tau_i P^{\text{OIS}}(t,t_i)} \tag{2.5}$$

Critically, $S^{3M}(t,t_0,t_n) \neq \frac{P^{\text{OIS}}(t,t_0)-P^{\text{OIS}}(t,t_n)}{A^{\text{OIS}}(t)}$ — the swap rate depends on the tenor-specific index curve, not the discount curve.

Remark 2.1. The multi-curve framework introduces basis risk between swaps with different floating leg tenors. A 3M LIBOR-based swap and a 6M LIBOR-based swap with identical fixed legs will not have the same par rate, even when using the same OIS discount curve.

2.3.3 Basis Spreads

Three types of basis spreads are fundamental. The LIBOR–OIS basis is the spread between LIBOR and OIS, reflecting credit risk and liquidity premia in the interbank market. It is incorporated via

$$L^{3M}(t,T,T+\delta) = L^{\text{OIS}}(t,T,T+\delta) + b_{\text{LO}}(t,T,\delta) \tag{2.6}$$

The tenor basis is the spread between different LIBOR tenors and reflects the term structure of credit and liquidity premia. A 3M-vs-6M basis swap exchanges 3M LIBOR for 6M LIBOR plus a spread $b_{3M,6M}$.

The cross-currency basis is the deviation from covered interest parity, observed particularly for JPY and EUR against USD during the crisis.

2.4 Collateralized Discounting

2.4.1 Framework

Following Piterbarg [59], the value of a collateralized derivative at time t paying an integrable payoff X_T at maturity is

Theorem 2.2 (Collateralized Derivative Pricing)**.** *Under perfect collateralization, with collateral remuneration rate $r^C(s)$ and with the collateral account defining an equivalent risk-neutral measure, the derivative value is*

$$V(t) = \mathbb{E}^{\mathbb{Q}}\left[\exp\left(-\int_t^T r^C(s)\,ds\right) X_T \mid \mathcal{F}_t\right] \tag{2.7}$$

Proof. The result follows from the collateralised self-financing replication argument in [59]. Under perfect collateralisation the funding account earns the collateral remuneration rate, and the value discounted by the collateral account is a true martingale by the standing admissibility assumptions. Taking its conditional expectation gives (2.7). □

Remark 2.3. For typical CSA agreements, $r^C(s)$ is the overnight rate (Fed Funds, EONIA, SONIA), implying OIS discounting. This resolves the apparent arbitrage from using LIBOR-based discounting for collateralized trades.

2.4.2 Impact on CMS Pricing

For a CMS payment $\tau S(T_{\text{fix}})$ at time T_{pay}, the value under OIS discounting is

$$V^{\text{CMS}}(t) = \tau P^{\text{OIS}}(t, T_{\text{pay}})\mathbb{E}^{T_{\text{pay}},C}[S(T_{\text{fix}})|\mathcal{F}_t] \tag{2.8}$$

where $\mathbb{E}^{T_{\text{pay}},C}$ denotes expectation under the collateralized T_{pay}-forward measure. The measure change from the annuity measure (natural for swap rates) to the collateralized forward measure introduces a convexity adjustment that differs from the classical uncollateralized case.

2.5 RFR-Based CMS Products

Following the LIBOR cessation (June 2023 for most major currencies), new CMS products are increasingly specified on risk-free rates (RFRs) or on fallback-adjusted swap rates rather than on legacy LIBOR fixings [60, 48, 51]. The compounded RFR over period $[T_{i-1}, T_i]$ is

$$C_i = \frac{1}{\tau_i}\left[\prod_{j:T_{i-1}<t_j\le T_i}(1 + r(t_{j-1})\delta_j) - 1\right], \tag{2.9}$$

where $r(t_j)$ is the overnight RFR rate.

Lemma 2.4 (Compounding Rate Equivalence)**.** *Assume the compounded overnight observation grid exactly spans $[T_{i-1}, T_i]$ and let $t \le T_{i-1}$. Under the T_i-forward measure,*

$$\mathbb{E}^{T_i}[C_i|\mathcal{F}_t] = L^{OIS}(t; T_{i-1}, T_i), \tag{2.10}$$

where L^{OIS} is the OIS forward rate.

Proof. In a collateralised OIS setting the compounded overnight payoff over $[T_{i-1}, T_i]$ satisfies

$$1 + \tau_i C_i = \frac{B^C(T_i)}{B^C(T_{i-1})},$$

where B^C is the collateral money-market account. Changing from the collateralised risk-neutral measure to the T_i-forward measure gives

$$\mathbb{E}^{T_i}[1 + \tau_i C_i \mid \mathcal{F}_t] = \frac{P^{\mathrm{OIS}}(t, T_{i-1})}{P^{\mathrm{OIS}}(t, T_i)}.$$

Rearranging yields

$$\mathbb{E}^{T_i}[C_i \mid \mathcal{F}_t] = \frac{1}{\tau_i}\left(\frac{P^{\mathrm{OIS}}(t, T_{i-1})}{P^{\mathrm{OIS}}(t, T_i)} - 1\right) = L^{\mathrm{OIS}}(t; T_{i-1}, T_i),$$

as in [34]. □

Remark 2.5. RFR-based CMS rates simplify to $S^{\mathrm{RFR}} = \frac{P^{\mathrm{OIS}}(t,T_0) - P^{\mathrm{OIS}}(t,T_n)}{A^{\mathrm{OIS}}}$ when the fixed and floating legs are consistently OIS based, eliminating the explicit LIBOR–OIS basis term. Under the exact-spanning, period-end payment assumptions of Lemma 2.4, the compounded RFR expectation equals the OIS forward. Additional timing, settlement or nonlinear-payoff features can introduce further convexity adjustments; for short accrual periods and small volatility these are often approximated by second-order terms, with more complete treatments given in [6, 71, 48].

2.6 Implications for CMS Pricing

The multi-curve framework has several implications. Convexity adjustments depend on correlation structure because the swap-rate drift under a payment-date forward measure contains covariations between OIS discount factors and tenor-specific index curves. Term-structure dynamics also matter: yield-curve tilts and curvature changes can materially affect CMS convexity, whereas models restricted to parallel shifts can underestimate the adjustment [29, 38, 50, 43]. Spread-contingent products require enough rate factors to generate independent movements of the relevant swap rates, and calibration is therefore high-dimensional when smile quotes across strikes and tenors must be matched while preserving realistic forward-rate correlations.

These challenges motivate the development of quasi-Gaussian HJM models presented in Chapter 4, which provide tractable closed-form or semi-analytic solutions while capturing the necessary term structure dynamics.

2.7 Summary

This chapter established the multi-curve framework used in the remainder of the thesis: OIS discounting for collateralised derivatives, tenor-specific index curves for forward-rate projection, basis spreads between LIBOR tenors and LIBOR–OIS curves, the collateralised derivative pricing identity in Theorem 2.2, RFR-based CMS products with compounding convexity, and the sensitivity of CMS convexity adjustments to term-structure tilts.

The next chapter presents the CMS product classes priced in the empirical study: CMS swaps, CMS spread swaps and options, callable CMS swaps, and callable CMS spread swaps.

3 CMS Products

This chapter fixes the contract notation used in the numerical comparison. The purpose is not to survey the full CMS product universe, but to specify precisely the payoffs that are priced later under the Gaussian and quasi-Gaussian model families. All cashflows are collateralised and discounted on the OIS curve of Chapter 2; forwarding curves enter only through the floating rate projections. The notation follows the CMS and multi-curve conventions in [29, 50, 33].

3.1 Forward Swap Rates

A swap rate is a term rate: it is the fixed rate on a swap whose fixed-leg present value equals the projected floating-leg present value.

Definition 3.1 (Multi-curve forward swap rate)**.** Let $T_a < T_{a+1} < \cdots < T_b$ be fixed-leg dates with accrual fractions τ_i, and let $P^d(t,T)$ denote the OIS discount factor. The annuity is

$$A^{a,b}(t) = \sum_{i=a+1}^{b} \tau_i P^d(t,T_i). \tag{3.1}$$

For a floating leg projected on a forwarding curve f, with floating accruals δ_i and forward rates $L^f_{i-1}(t)$ over $[T_{i-1},T_i]$, define the projected floating-leg present value by

$$\Pi^{a,b}_f(t) = \sum_{i=a+1}^{b} \delta_i P^d(t,T_i) L^f_{i-1}(t). \tag{3.2}$$

The multi-curve par swap rate is the fixed rate that makes the swap value zero

$$S^{a,b}(t) = \frac{\Pi^{a,b}_f(t)}{A^{a,b}(t)}. \tag{3.3}$$

Equivalently, if the floating and fixed schedules coincide and $L^d_{i-1}(t)$ denotes the simply compounded forward rate implied by the discount curve, then $\sum_i \delta_i P^d(t,T_i) L^d_{i-1}(t) = P^d(t,T_a) - P^d(t,T_b)$, so the par rate may be written as the single-curve par rate plus the discounted forwarding-basis cashflows; see [33, 27].

In the single-curve limit $L^f = L^d$, (3.3) reduces to the familiar expression $(P^d(t,T_a) - P^d(t,T_b))/A^{a,b}(t)$. A CMS rate is the random variable $S^{a,b}(T)$ fixed at a future date T and paid later, usually at the end of the coupon period.

3.2 CMS Coupons and CMS Swaps

Definition 3.2 (CMS coupon)**.** Fix a payment date T_i, fixing date T_{i-1}, notional N, accrual fraction τ_i, gearing α, and spread s. A CMS coupon linked to the underlying swap rate S^{a_i,b_i} pays

$$C_i^{\mathrm{CMS}} = N\tau_i\Big(\alpha S^{a_i,b_i}(T_{i-1}) + s\Big) \tag{3.4}$$

at T_i.

Definition 3.3 (CMS swap)**.** A CMS swap is a finite portfolio of CMS coupons $\{C_i^{\mathrm{CMS}}\}_{i=1}^n$ exchanged against a fixed or floating funding leg. Its time-zero value is

$$V_0^{\mathrm{CMS}} = \sum_{i=1}^{n} P^d(0,T_i)\,\mathbb{E}^{T_i}\Big[C_i^{\mathrm{CMS}}/N\Big]\,N - V_0^{\mathrm{funding}}, \tag{3.5}$$

where $\mathbb{Q}^{T_i}$ is the T_i-forward measure associated with the OIS discount bond numeraire.

The non-trivial term is $\mathbb{E}^{T_i}[S^{a_i,b_i}(T_{i-1})]$. The swap rate is a martingale under its annuity measure, not under the payment-date forward measure, so the expectation contains the CMS convexity adjustment derived in Chapter 5.

3.3 CMS Spread Coupons

Definition 3.4 (CMS spread coupon)**.** Let $S^{(1)}$ and $S^{(2)}$ be two CMS rates, typically a long-tail and a short-tail swap rate such as 10Y and 2Y. A CMS spread coupon pays

$$C_i^{\mathrm{spr}} = N\tau_i\,\Phi\Big(\alpha_1 S^{(1)}(T_{i-1}) - \alpha_2 S^{(2)}(T_{i-1}) + s\Big) \tag{3.6}$$

at T_i, where $\Phi(u) = u$ for a plain spread coupon, $\Phi(u) = u^+$ for a caplet-type payoff, and combinations of positive parts represent capped or floored coupons.

Pricing (3.6) requires the joint distribution of the two swap rates at the fixing date. This is the first point in the thesis where the factor structure of the interest-rate model becomes central: a one-factor model moves all swap rates through one state variable, whereas a multi-factor model can generate non-trivial spread volatility and term-structure correlation. CMS spread approximations and their dependence on spread correlation are discussed in [11, 50, 31].

3.4 Callable CMS Contracts

Definition 3.5 (Callable CMS product)**.** Let $0 < T_{e_1} < \cdots < T_{e_m}$ be issuer exercise dates and let $\mathcal{C}(t,u)$ denote the discounted value at time t of all contractual CMS

cashflows scheduled after t and up to u. The issuer may cancel the remaining cashflows at an exercise date. The investor value is the Bermudan optimal-stopping problem [57, 66]

$$V_0^{\text{call}} = \inf_{\tau \in \mathcal{T}} \mathbb{E}^{\mathbb{Q}} \left[\sum_{i:T_i \le \tau} D(0,T_i) C_i \right], \qquad \mathcal{T} \subset \{T_{e_1}, \ldots, T_{e_m}, T_n\}, \tag{3.7}$$

where $D(0,T_i)$ is the OIS stochastic discount factor and T_n denotes final maturity. Equivalently, the product is the non-callable CMS structure minus a Bermudan cancellation option held by the issuer.

The continuation value in (3.7) depends on the conditional distribution of future CMS rates. For this reason callable CMS products are model-dependent even when the non-callable CMS coupons can be priced by static replication. The Longstaff–Schwartz approximation used in the implementation is described in Chapter 6; see [46, 24] for regression Monte Carlo methods for such stopping problems.

The empirical chapters use the market shorthand $m\text{nc}NY$ for callable structures. Here N is the final maturity in years and $m$nc denotes an initial m-year non-call period. Thus 1nc10Y is a ten-year product that is first callable after one year, whereas 5nc10Y is a ten-year product with a five-year non-call period.

3.5 Market Conventions

The empirical study uses EUR conventions: OIS/ESTR discounting, EURIBOR 6M forwarding, annual fixed-leg coupons with 30/360 day count, Act/360 floating accrual, TARGET calendar, and Modified Following business-day adjustment. These conventions determine the cashflow dates and accrual fractions in the formulas above. They do not change the mathematical structure of the comparison, but they are required for reproducing the numerical results of Chapter 7.

3.6 RFR-Based CMS Products

In an RFR setting the forward-looking IBOR rate in the floating leg is replaced by a backward-looking compounded overnight rate. The formal definitions above remain valid after replacing L^f with the compounded RFR projection over the coupon period. The convexity adjustment then contains an additional compounding-convexity component, which is typically small for the EUR examples considered here and is not the focus of the empirical comparison; see [34] for RFR and overnight-index

conventions.

3.7 Summary

The thesis prices four contract classes: CMS swaps, CMS spread swaps, callable CMS swaps, and callable CMS spread swaps. Non-callable products require convexity-adjusted expectations of one or two swap rates. Callable products add an optimal stopping problem and therefore depend on the model-implied joint dynamics of the relevant swap rates. This separation explains why plain CMS products are useful controls in Chapter 7, while callable CMS spread products are the main test of the multi-factor QG-SV model.

4 Gaussian and Quasi-Gaussian Models

This chapter introduces the four model families used in the empirical comparison of Chapter 7: a one-factor Gaussian short-rate model (LGM/Hull–White), a two-factor Gaussian short-rate model (G2++), the one-factor quasi-Gaussian family (with local and stochastic volatility), and the multi-factor quasi-Gaussian model with stochastic volatility (MF QG-SV). The QG family belongs to the Cheyette/finite-dimensional HJM lineage [17, 64, 15, 8, 9]. The chapter is organised as a model nesting hierarchy: each model is a parameter restriction of the next more general one. The thesis contribution is not a new QG-SV specification, but the numerical adaptation of the Andersen–Piterbarg benchmark-rate parametrisation, together with the piecewise MC bootstrap workflow used later in the empirical study.

For the standard Gaussian and quasi-Gaussian results we follow Andersen and Piterbarg [5], Brigo and Mercurio [11], and the earlier Cheyette finite-dimensional HJM construction [17]. Background on HJM and on the quasi-Gaussian construction is summarised here only to the extent needed for Chapter 7; the main derivations used later in the thesis are collected in Appendix A.

4.1 Probability-Theoretic Setup

We work on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{Q})$ supporting a d-dimensional Brownian motion $\mathbf{W}^x$ for the rate factors and an independent Brownian motion W^z for the variance factor. The filtration is the augmented natural filtration. We adopt the multi-curve setting of Chapter 2: one curve $P^d(\cdot,\cdot)$ for collateralised discounting and a forwarding curve $P^f(\cdot,\cdot)$ per relevant tenor, with both built consistently from market quotes. Standard regularity conditions on volatility coefficients (Borel measurability in time, local Lipschitz continuity in the state, and at most linear growth) are assumed so that the SDEs below have unique strong solutions.

4.2 The HJM Viewpoint and Finite-Dimensional Realisations

The models used in this thesis are best understood as finite-dimensional representations of an HJM term-structure model. In the HJM framework one models the entire instantaneous forward-rate curve $T \mapsto f(t,T)$, not only the short rate $r(t) = f(t,t)$.

The OIS discount bond is then

$$P^d(t,T) = \exp\left(-\int_t^T f(t,u)\,du\right), \tag{4.1}$$

and, under the collateralised risk-neutral measure, a generic forward-rate dynamics has the form

$$df(t,T) = \alpha_f(t,T)\,dt + \boldsymbol{\sigma}_f(t,T)^\top d\mathbf{W}_t. \tag{4.2}$$

The central HJM insight is that the drift is not a free modelling input. Absence of arbitrage fixes it through the volatility structure

$$\alpha_f(t,T) = \boldsymbol{\sigma}_f(t,T)^\top \int_t^T \boldsymbol{\sigma}_f(t,u)\,du. \tag{4.3}$$

Equivalently, once the forward-curve volatility is chosen, the bond-price drift is forced to be $r(t)$ under the risk-neutral measure. This is why the models below focus on volatility, factor loading and covariance structure: the drift adjustments are determined by no-arbitrage rather than estimated independently [32, 5].

The price of this generality is that an unrestricted HJM model is infinite-dimensional. A simulation path must in principle carry the whole forward curve $f(t,\cdot)$. A finite-dimensional realisation is a restriction of the HJM volatility such that the evolving curve remains a deterministic function of finitely many Markov state variables. The Cheyette and Ritchken–Sankarasubramanian constructions, the later geometric finite-dimensional-realisation results, and recent multi-curve extensions identify when such reductions are possible [17, 64, 15, 8, 9, 20]. The quasi-Gaussian models in this thesis belong to that class.

The intuition is simple. Mean reversion gives each Brownian factor a maturity loading: a shock at time t affects the T-maturity bond through the integrated loading

$$G_k(t,T) = \int_t^T \exp\left(-\int_t^u \kappa_k(s)\,ds\right) du. \tag{4.4}$$

The state variable $x_k(t)$ records the accumulated random displacement of the curve in this loading direction. The auxiliary variables $y_{kl}(t)$ record the accumulated covariance corrections required by the HJM drift restriction. Together they make the whole discount curve a function of the finite state

$$P^d(t,T) = \frac{P^d(0,T)}{P^d(0,t)} \exp\left(-\sum_k G_k(t,T)x_k(t) - \frac{1}{2}\sum_{k,l} G_k(t,T)G_l(t,T)y_{kl}(t)\right). \tag{4.5}$$

Thus x describes how the curve moves, while y enforces the no-arbitrage convexity correction induced by those moves. In the Gaussian case y is deterministic; in local-volatility and stochastic-volatility QG models it becomes path-dependent because the instantaneous covariance depends on the state and, in QG-SV, on the CIR variance factor $z(t)$.

This perspective also explains why QG models are natural for CMS products. CMS convexity and callable exercise depend on how swap rates move across future dates, not only on today's curve. The HJM volatility loading $G_k(t,T)$ determines how each factor changes different maturities, and therefore how level, slope and curvature risk enter CMS spread volatility. The QG-SV model keeps this arbitrage-free HJM curve structure while adding displacement and stochastic volatility to fit the observed swaption smile.

4.3 One-Factor Gaussian Short Rate Model

The one-factor Gaussian short-rate model (LGM/Hull–White) defines a single Markov state $x(t)$ with mean reversion κ and time-dependent volatility $\sigma(t)$. It is the time-inhomogeneous Gaussian short-rate benchmark descended from the Vasicek, Ho–Lee and Hull–White models [70, 36, 37, 11]. In the Cheyette/HJM coordinates used by the QG formulas below, the Gaussian restriction is written as

$$\begin{aligned} dx(t) &= \Big(y(t) - \kappa x(t)\Big)\,dt + \sigma(t)\,dW(t), \\ dy(t) &= \Big(\sigma^2(t) - 2\kappa y(t)\Big)\,dt, \qquad x(0) = y(0) = 0, \end{aligned} \tag{4.6}$$

and the short rate is $r(t) = f(0,t) + x(t)$. The auxiliary variance $y(t)$ is deterministic in this Gaussian case and is exactly the drift correction that preserves the initial term structure under the bond representation. Equivalently, after deterministic re-centring, the same model can be written in the standard shifted Hull–White form $r(t) = \varphi(t) + \tilde{x}(t)$, where $d\tilde{x} = -\kappa\tilde{x}\,dt + \sigma(t)dW(t)$ and φ is chosen to fit the initial discount curve. Bond prices are exponential-affine,

$$P(t,T) = \frac{P(0,T)}{P(0,t)} \exp\Big(-G(t,T)\,x(t) - \tfrac{1}{2}G(t,T)^2 y(t)\Big), \tag{4.7}$$

where

$$G(t,T) = \frac{1 - e^{-\kappa(T-t)}}{\kappa}, \qquad y(t) = \int_0^t e^{-2\kappa(t-u)}\sigma^2(u)\,du.$$

European swaptions admit a closed-form Black/Bachelier price after Jamshidian decomposition [39], which is exploited by the LGM benchmark in Chapter 7. The

one-factor model produces a flat smile at every expiry: this is its principal limitation and the motivation for moving to QG-SV when calibrating to a smile surface.

4.4 Multi-Factor Gaussian Model (G2++)

The two-factor Gaussian model G2++ extends the one-factor shifted Gaussian short-rate model with two correlated mean-reverting factors [11]

$$\begin{aligned} dx_1(t) &= -\kappa_1 x_1(t)\,dt + \sigma_1(t)\,dW_1(t), \\ dx_2(t) &= -\kappa_2 x_2(t)\,dt + \sigma_2(t)\,dW_2(t), \qquad d\langle W_1, W_2\rangle_t = \rho\,dt, \\ r(t) &= \varphi(t) + x_1(t) + x_2(t). \end{aligned} \tag{4.8}$$

Here the deterministic shift $\varphi(t)$ is chosen so that the model reproduces the initial discount curve exactly. Bond prices remain exponential-affine in the state, with the deterministic affine prefactor absorbing the initial curve and variance adjustment, and European swaptions can be priced semi-analytically via the standard Brigo–Mercurio Gaussian integral [11]. The additional factor introduces a controlled non-trivial yield-curve covariance structure: short-end and long-end rates can move differently, which is important for callable products whose payoffs depend on a yield-curve slope. As with the one-factor case, G2++ produces a flat smile at every expiry. We use G2++ as the multi-factor Gaussian benchmark in Section 7.2.1.

4.5 One-Factor Quasi-Gaussian Model with Local Volatility

The one-factor quasi-Gaussian construction can be viewed as the Cheyette finite-dimensional HJM model written in the notation of Andersen and Piterbarg [17, 5]. It extends the Gaussian short-rate model by allowing the diffusion coefficient to depend on the state. With one factor, the dynamics are

$$dx(t) = (y(t) - \kappa\, x(t))\,dt + \sigma_r(t,x,y)\,dW(t), \qquad dy(t) = (\sigma_r^2(t,x,y) - 2\kappa\, y(t))\,dt, \tag{4.9}$$

where $y(t)$ is a deterministic-in-state-but-pathwise-stochastic auxiliary variance that absorbs the path-dependence of the HJM forward-rate evolution. In the Andersen–Piterbarg notation [5, Chapter 13], this Markovianisation gives the following exponential bond representation.

Theorem 4.1 (Quasi-Gaussian bond representation)**.** *Under the standing regularity assumptions on σ_r, the time-t price of a zero-coupon bond maturing at T is determined*

by the two Markov state variables $(x(t), y(t))$ *through*

$$P(t,T) = \frac{P(0,T)}{P(0,t)} \exp\Big(-G(t,T)\,x(t) - \tfrac{1}{2}G(t,T)^2\,y(t)\Big). \tag{4.10}$$

The state-dependent diffusion $\sigma_r(t,x,y)$ generates a smile because the option payoff samples the volatility at non-zero x. A common parametric choice is displaced diffusion, $\sigma_r(t,x,y) = \lambda(t)\Big((1-b(t))\,f(0,t) + b(t)\,(f(0,t)+x)\Big)$, with $b(t) \in [0,1]$ controlling the skew [28, 5]. Local-volatility QG produces a static smile that matches the spot market well but is rigid in the forward direction; this limitation is analogous to the motivation for stochastic-volatility extensions in LIBOR and swap-rate market models [2, 58, 42].

4.6 One-Factor QG-SV Model

The one-factor QG-SV model adds a CIR variance factor $z(t)$ that scales the rate diffusion, combining the CIR/Heston square-root variance mechanism with the quasi-Gaussian interest-rate state [18, 35, 58, 5]

$$\begin{aligned} dx(t) &= (y(t) - \kappa x(t))\,dt + \sqrt{z(t)}\,\sigma_r(t,x,y)\,dW^x(t), \\ dz(t) &= \theta(\bar{z} - z(t))\,dt + \eta\sqrt{z(t)}\,dW^z(t). \end{aligned} \tag{4.11}$$

with $z(0) = \bar{z} = 1$ in the numerical specification (variance is normalised so that σ_r is the level and z is its multiplicative perturbation). The auxiliary variance dynamics become $dy(t) = (z(t)\sigma_r^2(t,x,y) - 2\kappa\,y(t))\,dt$, which generalises the local-volatility case continuously through $z \equiv 1$.

Remark 4.2 (Design choice $\rho_{xz} = 0$)**.** We deliberately take the rate–variance correlation to be zero, $\langle dW^x, dW^z\rangle_t = 0$. This is consistent with the Heston/CIR analytical machinery in [35, 5], preserves the CIR form of z under the annuity-measure changes used here, and is sufficient on the EUR data set of Chapter 7 to fit the swaption smile to low-single-digit basis-point residuals. A non-zero ρ_{xz} would add one more weakly identified parameter; estimating whether it improves fit robustly is outside the scope of the case study. Skew is generated by the displacement parameter $b(t)$, not by ρ_{xz}.

European swaption pricing in QG-SV requires either Monte Carlo or a semi-analytical Fourier formula obtained by freezing the swap-rate sensitivity at its time-zero value:

Remark 4.3 (Semi-analytical approximation)**.** A semi-analytical (SA) pricer can be built in the QG-SV model by freezing the swap-rate sensitivity $\zeta(t, x, y)$ at $t = 0$ and applying a displaced-Heston Fourier inversion [35, 13, 5]. This is fast (millisecond per swaption), but the freezing step removes the positive co-movement between $z(t)$ and $\zeta(t, x, y)^2$, biasing the implied variance. The sign and order of the bias are made precise by the following proposition.

Proposition 4.4 (Sufficient condition for SA–MC variance bias)**.** *Let the exact annuity-measure swap-rate dynamics in Theorem 4.10 be written in scalar form as*

$$dS_t = \sqrt{z_t}\,\zeta(t, X_t)\,dW_t^A,$$

where $X_t = (\mathbf{x}(t), \mathbf{y}(t))$*. Let the semi-analytical approximation replace* $\zeta(t, X_t)$ *by a deterministic frozen sensitivity* $\zeta_0(t)$*. Define* $V_{\mathrm{MC}}(T) = \langle S \rangle_T$ *and* $V_{\mathrm{SA}}(T) = \int_0^T z_t \zeta_0(t)^2 dt$*. Then the expected integrated variance difference is*

$$\Delta V_T = \int_0^T \mathbb{E}^A\Big[z_t\{\zeta(t, X_t)^2 - \zeta_0(t)^2\}\Big]\,dt. \tag{4.12}$$

In particular, if $\mathbb{E}^A[z_t] > 0$*,* $\zeta_0(t)^2 \le \mathbb{E}^A[\zeta(t, X_t)^2]$ *and* $\mathrm{Cov}(z_t, \zeta(t, X_t)^2) \ge 0$ *for* $0 \le t \le T$*, then the frozen approximation underestimates the exact integrated variance.*

Proof. By Theorem 4.10, the instantaneous quadratic variation of the swap rate is $d\langle S \rangle_t = z_t \zeta(t, X_t)^2 dt$. Hence the pathwise integrated variance equals

$$V_{\mathrm{MC}}(T) = \int_0^T z_t\,\zeta(t, X_t)^2\,dt. \tag{4.13}$$

The frozen approximation gives

$$V_{\mathrm{SA}}(T) = \int_0^T z_t\,\zeta_0(t)^2\,dt.$$

Subtracting expectations and applying Fubini's theorem gives (4.12). The pointwise integrand can be written as

$$\mathbb{E}^A[z_t\{\zeta(t, X_t)^2 - \zeta_0(t)^2\}] = \mathbb{E}^A[z_t]\Big(\mathbb{E}^A[\zeta(t, X_t)^2] - \zeta_0(t)^2\Big) + \mathrm{Cov}(z_t, \zeta(t, X_t)^2).$$

The stated sufficient condition and $\mathbb{E}^A[z_t] > 0$ make both terms on the right-hand side non-negative. □

Proposition 4.4 explains the empirical SA–MC inconsistency discussed in Sec-

tion 7.4: SA-calibrated parameters absorb a missing covariance term into the fitted volatility parameters, and those parameters need not remain appropriate when transferred into the full Monte Carlo dynamics. Appendix A gives the corresponding integrated-variance derivation. For this reason the callable-pricing results in this thesis use Monte Carlo calibration when the downstream pricing method is Monte Carlo.

4.7 Multi-Factor QG-SV Model

The multi-factor QG-SV model is the central object of this thesis. It generalises the Cheyette/QG finite-dimensional HJM construction to d correlated rate factors driven by a single CIR variance factor, retains the exponential-affine bond price in the rate state, and admits a piecewise parameter representation that supports a sequential bootstrap calibration. Most of the standard textbook material on the multi-factor QG model with deterministic volatility in Andersen–Piterbarg [5, §13.1] carries over without change: the only structural addition is the variance factor $z(t)$.

Definition 4.5 (Multi-factor QG-SV; [5], Chapter 13)**.** The model has d rate factors $\mathbf{x}(t) = (x_1(t), \ldots, x_d(t))$, the symmetric cross-variance matrix $\mathbf{y}(t) = (y_{kl}(t))_{k,l=1}^{d}$, and a single CIR variance factor $z(t)$. Under the risk-neutral measure $\mathbb{Q}$,

$$dx_k(t) = \left(\sum_{l=1}^{d} y_{kl}(t) - \kappa_k\, x_k(t)\right) dt + [\sigma_x(t, \mathbf{x}, \mathbf{y})]_{k,\cdot}^{\top} \sqrt{z(t)}\, d\mathbf{W}^x(t), \tag{4.14}$$

$$dy_{kl}(t) = \left(z(t)\, [\sigma_x \sigma_x^{\top}]_{kl} - (\kappa_k + \kappa_l)\, y_{kl}(t)\right) dt, \tag{4.15}$$

$$dz(t) = \theta(\bar{z} - z(t))\, dt + \eta \sqrt{z(t)}\, dW^z(t), \tag{4.16}$$

with $\langle dW_i^x, dW_j^x \rangle_t = \delta_{ij}\, dt$ and $\langle d\mathbf{W}^x, dW^z \rangle_t = 0$ as in Remark 4.2. Inter-factor benchmark-rate correlation is embedded in the diffusion matrix of Definition 4.6.

4.7.1 Benchmark-Rate Parametrisation

For the empirical comparison we use the benchmark-rate displaced-diffusion parametrisation of Andersen–Piterbarg [5, §13.1.2–13.1.4], which expresses the diffusion of $\mathbf{x}$ in terms of the volatilities of d benchmark forward rates rather than directly in state space.

Definition 4.6 (Benchmark-rate MF QG-SV parametrisation)**.** Each rate factor $k = 1, \ldots, d$ is associated with a benchmark tenor Δ_k. The k-th benchmark forward

rate is

$$f_k(t) = f(0, t+\Delta_k) + \sum_{j=1}^{d} e^{-\kappa_j \Delta_k} \Big(x_j(t) + \sum_{l=1}^{d} y_{jl}(t)\, G_l(t, t+\Delta_k) \Big), \tag{4.17}$$

and its volatility is given by the displaced-diffusion form

$$\sigma_k^f(t) = \lambda_k(t) \Big((1 - b_k(t))\, f(0, t+\Delta_k) + b_k(t)\, f_k(t) \Big), \tag{4.18}$$

with $b_k = 0$ giving normal dynamics and $b_k = 1$ giving log-normal dynamics for the k-th benchmark rate. The state diffusion is then constructed by the linear transformation

$$\sigma_x(t, \mathbf{x}, \mathbf{y}) = H(t)\, (H^f)^{-1}\, \mathrm{diag}(\sigma_1^f, \dots, \sigma_d^f)\, D^f, \tag{4.19}$$

where $H_{ik}^f = e^{-\kappa_k \Delta_i}$, $H(t) = \mathrm{diag}(e^{-\kappa_1 t}, \dots, e^{-\kappa_d t})$, and D^f is the lower-triangular Cholesky factor of the inter-factor correlation matrix $\rho_{kl} = \exp(-\alpha |\Delta_k - \Delta_l|)$. The variance follows the CIR dynamics of Definition 4.5.

The structural parameters $(\kappa_k, \Delta_k, \alpha, \theta, \bar{z})$ are fixed before calibration; the calibrated parameters $(\lambda_k(t), b_k(t), \eta(t))$ are time-dependent and represented as piecewise constants. Two routine but important properties of this parametrisation are worth stating explicitly.

Remark 4.7 (State-space dimension)**.** The total state vector $(\mathbf{x}, \mathbf{y}, z)$ has dimension

$$d + \tfrac{d(d+1)}{2} + 1 \;=\; \tfrac{d(d+3)}{2} + 1, \tag{4.20}$$

so the $d = 2$ specification carries six state variables $(x_1, x_2, y_{11}, y_{22}, y_{12}, z)$ and the $d = 4$ variant used in the dimensionality study of Section 7.3.8 carries fifteen.

Remark 4.8 (Feller condition for the variance factor)**.** The CIR variance factor $z(t)$ stays strictly positive provided $2\theta\bar{z} \geq \eta^2$ (Feller condition) [19, 18]. The calibrated $(\theta, \bar{z}, \eta)$ in the empirical study of Chapter 7 do not always satisfy the strict inequality at every expiry; the numerical scheme therefore uses a moment-matched Gaussian approximation to the CIR transition and floors negative draws at zero. See Chapter 6 for the discretisation.

Theorem 4.9 (MF QG-SV bond pricing)**.** *The bond price is*

$$P(t,T) = \frac{P(0,T)}{P(0,t)} \exp\Big(-\sum_{k=1}^{d} G_k(t,T)\, x_k(t) - \tfrac{1}{2} \sum_{k,l=1}^{d} G_k(t,T)\, G_l(t,T)\, y_{kl}(t) \Big), \tag{4.21}$$

with $G_k(t,T) = \int_t^T e^{-\int_t^u \kappa_k(s)\,ds}\,du$. *The variance factor* $z(t)$ *enters only through the dynamics of* $(\mathbf{x},\mathbf{y})$ *and not through the functional form, which is what allows* z *to be a single factor shared across all* d *rate dimensions while preserving tractable swap-rate dynamics in the annuity measure.*

Theorem 4.10 (Swap rate dynamics under the annuity measure)**.** *Let* $A^{a,b}(t) = \sum_{i=a+1}^{b} \tau_i\, P^d(t,T_i)$ *be the annuity of a fixed leg with payment dates* $T_a < \cdots < T_b$*, and let* $S^{a,b}(t)$ *denote the OIS-equivalent forward par swap rate* $(P^d(t,T_a) - P^d(t,T_b))/A^{a,b}(t)$*. Under Assumption 0.4 and the annuity measure* $\mathbb{Q}^A$ *with numeraire* $A^{a,b}$*,* $S^{a,b}$ *is a true martingale on the finite swap horizon and has diffusion*

$$\begin{aligned} dS^{a,b}(t) = \sum_{k=1}^{d} \nabla_{x_k} S^{a,b}(t)\,\big[\sigma_x(t,\mathbf{x},\mathbf{y})\big]_{k,\cdot} \\ \times \sqrt{z(t)}\,d\mathbf{W}^A(t). \end{aligned} \tag{4.22}$$

where $\mathbf{W}^A$ *is a* d*-dimensional Brownian motion under* $\mathbb{Q}^A$ *with independent components, and the swap-rate gradient with respect to* x_k *is*

$$\begin{aligned} \nabla_{x_k} S^{a,b}(t) = -\frac{1}{A^{a,b}(t)}\Big(& G_k(t,T_a)P^d(t,T_a) - G_k(t,T_b)P^d(t,T_b) \\ & - S^{a,b}(t)\sum_{i=a+1}^{b} \tau_i G_k(t,T_i)P^d(t,T_i)\Big). \end{aligned} \tag{4.23}$$

For a multi-curve par rate with deterministic projected-basis cashflows, the same annuity-measure martingale argument applies after replacing the displayed gradient by the quotient derivative of the projected floating-leg present value divided by $A^{a,b}$*. The formula above is therefore the zero-basis special case, not a general stochastic-basis identity.*

Proof. Itô's lemma applied to $S^{a,b} = (P^d(\cdot,T_a) - P^d(\cdot,T_b))/A^{a,b}$ together with the bond pricing formula of Theorem 4.9 yields the stated diffusion. The numerator is the value of a traded bond portfolio and $A^{a,b}$ is the numeraire, so their ratio is a $\mathbb{Q}^A$-martingale by the change-of-numeraire theorem and Assumption 0.4; hence the drift is zero. The factor $\sqrt{z(t)}$ enters multiplicatively because z scales the diffusion of every x_k through Definition 4.5. □

The instantaneous quadratic variation of the swap rate is

$$\frac{d\langle S^{a,b}\rangle_t}{dt} \;=\; z(t)\,\nabla_x S^{a,b}(t)^\top \sigma_x(t,\mathbf{x},\mathbf{y})\sigma_x(t,\mathbf{x},\mathbf{y})^\top \nabla_x S^{a,b}(t), \tag{4.24}$$

which makes explicit the three pricing-relevant ingredients of the QG-SV model: the rate-factor sensitivities $\nabla_{x_k} S$ that reflect the shape of the discount curve, the benchmark-rate correlation embedded in $\sigma_x \sigma_x^\top$, and the variance factor $z(t)$ that drives the smile.

Proposition 4.11 (Locality of calibrated buckets)**.** *Under the continuous-time piecewise parametrisation of Definition 4.12, with fixed structural functions κ_k and fixed tenor map Δ_k, the model price of a European swaption expiring at T_n depends on the calibrated buckets $\theta^{(1)}, \dots, \theta^{(n)}$ only. Later calibrated buckets $\theta^{(j)} = (\lambda^{(j)}, b^{(j)}, \eta^{(j)})$, $j > n$, cannot change the exact continuous-time value of the earlier-expiring instrument. If structural functions that enter the post-expiry bond loadings, such as κ_k, were also calibrated in later buckets, this locality would no longer follow. The sequential bootstrap used below is therefore a well-defined forward-induction calibration convention for the stated parametrisation, not a claim of equality with the fully unconstrained global optimum or of exact finite-path optimiser convergence.*

Proof. Both $(\mathbf{x}(t), \mathbf{y}(t))$ and $z(t)$ are Itô diffusions whose coefficients on $[T_{j-1}, T_j)$ depend only on $\theta^{(j)}$ by Definition 4.12. The strong Markov property propagates this dependence forward in time, so the stopped state at T_n is a measurable functional of $(\theta^{(1)}, \dots, \theta^{(n)})$ and the driving Brownian path on $[0, T_n]$. The payoff of a physically settled swaption is then obtained by applying a terminal swap-value map to that state. By the bond representation in Theorem 4.9, that map contains the deterministic loadings $G_k(T_n, T_i)$, which are fixed by κ_k, the initial curve and the payment schedule, not by later calibrated (λ, b, η) buckets. Therefore the expectation defining the T_n-expiring swaption price depends only on the first n calibrated buckets. The caveat about calibrating structural functions follows from the same formula, since changing post-expiry κ_k would change $G_k(T_n, T_i)$ for later payment dates. □

4.7.2 Piecewise Parametrisation and Bootstrap Calibration

Definition 4.12 (Piecewise parameters)**.** Given calibration expiries $0 = T_0 < T_1 < \cdots < T_N$, the calibrated parameters are constant on each interval, $\lambda_k(t) = \lambda_k^{(n)}$, $b_k(t) = b_k^{(n)}$, $\eta(t) = \eta^{(n)}$ for $t \in [T_{n-1}, T_n)$. There are $2d + 1$ free parameters per interval and $N(2d + 1)$ in total.

With fixed structural loadings, this piecewise structure decouples the continuous-time calibration problem in time: the model price of a swaption expiring at T_n depends on $(\theta^{(1)}, \dots, \theta^{(n)})$ only, which permits sequential bootstrap. Algorithm 1

states the procedure used in Chapter 6; the optimiser, weights, and Monte Carlo settings are deferred to that chapter.

Algorithm 1 MC bootstrap calibration of MF QG-SV

Input: swaption smile $\{(T_n, K_{n,j}, \sigma_n^{\mathrm{mkt}}(K_{n,j}))\}_{n=1}^{N}$
Input: structural parameters $(\kappa_k, \Delta_k, \alpha, \theta, \bar{z})$
Input: Monte Carlo settings (Sobol paths, time discretisation, antithetic flag)
for $n = 1$ to N **do**
    freeze $(\theta^{(1)*}, \ldots, \theta^{(n-1)*})$
    optimise $\theta^{(n)} = (\lambda_k^{(n)}, b_k^{(n)}, \eta^{(n)})$ over strikes $\{K_{n,j}\}$
    store $\theta^{(n)*}$ and the per-expiry RMS
end for
Output: piecewise parameter set $\{\theta^{(n)*}\}_{n=1}^{N}$

4.7.3 Progressive Model Complexity by Parameter Restrictions

Proposition 4.13 (Model nesting)**.** *The MF QG-SV model nests the multi-factor Gaussian and the multi-factor QG with skew as parameter restrictions. The restriction* $b_k \equiv 0$ *together with* $\eta \to 0^+$ *gives the multi-factor Gaussian model with no smile. The restriction* $\eta \to 0^+$ *with* b_k *free gives the multi-factor QG model with displaced-diffusion skew but no stochastic-volatility curvature. With no such restrictions, the* $2d + 1$ *interval parameters define the full QG-SV model.*

This nesting underpins the progressive-complexity comparison of Chapter 7 (Table 7.9): the same calibration procedure, the same Monte Carlo paths, and the same Longstaff–Schwartz callable pricing method are used in all three variants, with only the admissible parameter set changing between runs. Differences in callable NPV are therefore primarily interpreted as model-degree-of-freedom effects, subject to the residual calibration, discretisation and regression errors documented in Chapter 6.

4.8 Comparative Properties

The four families used in Chapter 7 differ along two axes: yield-curve covariance structure (number of rate factors) and smile structure (deterministic vs. stochastic volatility).

Within this comparison set, the QG-SV family is the only specification designed to reproduce the EUR ATM curve, smile skew, and smile curvature simultaneously. Whether this added flexibility is economically material on callable CMS-spread products is the empirical question studied in Chapter 7.

Table 4.1: Model features used in the empirical comparison.

Model	Rate factors	Smile	Forward smile	Pricing of vanillas
LGM (1F Gaussian)	1	flat	flat	closed form
G2++ (2F Gaussian)	2	flat	flat	semi-analytic
QG with skew	1–d	static skew	rigid	Monte Carlo
QG-SV (full)	1–d	full smile	moves with z	Monte Carlo

4.9 Summary

The four model families enter the comparison as a nested hierarchy. LGM and G2++ supply Gaussian benchmarks with closed-form European pricing. The one-factor QG and QG-SV models introduce smile via state-dependent and stochastic-volatility diffusions. The multi-factor QG-SV model of Section 4.7, in the benchmark-rate displaced-diffusion parametrisation of Definition 4.6, is the central model of this thesis. Its piecewise parameter structure (Definition 4.12) supports the sequential MC bootstrap (Algorithm 1) and the progressive-complexity nesting (Proposition 4.13) that drive the empirical comparison in Chapter 7. The next chapter develops the CMS pricing techniques required to evaluate these models on the product universe of Chapter 3.

5 CMS Pricing

This chapter derives the pricing identities used to value the products of Chapter 3. The main point is measure mismatch. A CMS coupon pays a swap rate at a payment date, but the swap rate is a martingale under its annuity measure, not under the payment-date forward measure. The CMS convexity adjustment is exactly the correction induced by this change of numeraire. Callable products then add an optimal stopping problem. The change-of-numeraire, CMS replication and stopping arguments follow the standard treatments in [22, 66, 11, 4, 29, 38, 56, 57].

5.1 Convexity Adjustment

Proposition 5.1 (CMS convexity adjustment)**.** *Consider a CMS coupon paying $\tau S^{a,b}(T)$ at $T_p \geq T$, where $S^{a,b}$ is the swap rate of Definition 3.1 and $A^{a,b}$ is the corresponding annuity. Assume Assumption 0.4 holds for the coupon payoff and for the annuity and payment-date numeraires, and that $S^{a,b}$ is a true martingale under the annuity measure. Its time-zero value is*

$$V_0 = \tau P^d(0,T_p)\mathbb{E}^{T_p}\left[S^{a,b}(T)\right]. \tag{5.1}$$

Let $\mathbb{Q}^A$ denote the annuity measure with numeraire $A^{a,b}$. Then

$$\mathbb{E}^{T_p}\left[S^{a,b}(T)\right] = \mathbb{E}^A\left[S^{a,b}(T)\frac{P^d(T,T_p)A^{a,b}(0)}{A^{a,b}(T)P^d(0,T_p)}\right], \tag{5.2}$$

and the convexity adjustment is

$$\mathrm{CA}(T,T_p) = \mathbb{E}^{T_p}\left[S^{a,b}(T)\right] - S^{a,b}(0). \tag{5.3}$$

Proof. The first valuation formula follows from the payment-date forward measure, whose density process is well defined by Assumptions 0.1 and 0.4. The Radon–Nikodym derivative between the payment-date forward measure and the annuity measure at time T is

$$\left.\frac{d\mathbb{Q}^{T_p}}{d\mathbb{Q}^A}\right|_{\mathcal{F}_T} = \frac{P^d(T,T_p)}{A^{a,b}(T)}\frac{A^{a,b}(0)}{P^d(0,T_p)}. \tag{5.4}$$

The density is positive and has unit expectation under $\mathbb{Q}^A$. Multiplying $S^{a,b}(T)$ by this density gives (5.2). Since $S^{a,b}$ is a martingale under $\mathbb{Q}^A$, $\mathbb{E}^A[S^{a,b}(T)] = S^{a,b}(0)$, and the difference between the two expectations is the convexity adjustment. □

Equation (5.2) is exact. Practical CMS formulas differ in how they approximate the random density $P^d(T,T_p)/A^{a,b}(T)$ as a function of the swap rate. The swap-yield mapping and Hagan's linear terminal swap-rate mapping are two standard choices [29, 38, 56, 61, 16, 4]. In the empirical study, plain CMS coupons are priced by analytical replication, while callable products are valued by Monte Carlo under the selected model.

5.2 Static Replication of CMS Payoffs

The following is the standard Carr–Madan static-replication identity, written under the swap annuity measure and with swaptions as the traded option instruments [14, 38].

Theorem 5.2 (Static replication under the annuity measure)**.** *Let $S_T = S^{a,b}(T)$ be an integrable swap rate that is a true martingale under its annuity measure, and let $F = S^{a,b}(0)$. Suppose $h : \mathbb{R} \to \mathbb{R}$ is twice continuously differentiable and*

$$\int_F^\infty |h''(K)|\, C(K)\, dK + \int_{-\infty}^F |h''(K)|\, P(K)\, dK < \infty.$$

Then

$$\mathbb{E}^A[h(S_T)] = h(F) + \int_F^\infty h''(K)\, C(K)\, dK + \int_{-\infty}^F h''(K)\, P(K)\, dK, \tag{5.5}$$

where $C(K) = \mathbb{E}^A[(S_T - K)^+]$ and $P(K) = \mathbb{E}^A[(K - S_T)^+]$ are payer and receiver swaption prices in annuity units.

Proof. For each realization s, the second-order Taylor representation with integral remainder gives

$$h(s) = h(F) + h'(F)(s-F) + \int_F^\infty h''(K)(s-K)^+\, dK + \int_{-\infty}^F h''(K)(K-s)^+\, dK. \tag{5.6}$$

Taking $\mathbb{Q}^A$-expectations is justified by the integrability condition and Fubini's theorem. Since $\mathbb{E}^A[S_T - F] = 0$, this yields (5.5). □

For a CMS coupon, the formal one-dimensional reduction is not simply $h(s) = s$. It is $h(s) = s\alpha_T(s)$, where α_T is the conditional annuity mapping in Proposition 5.3. This representation is exact only at the level of the conditional mapping; practical CMS formulas replace α_T by a tractable approximation, for example a swap-yield or linear terminal-swap-rate mapping. With this approximation fixed, the market

swaption smile enters through $C(K)$ and $P(K)$ across strikes, not only through the ATM volatility. The same static-replication logic applies to caplets, floorlets, and capped/floored CMS coupons by choosing the corresponding payoff function h.

Proposition 5.3 (One-dimensional CMS replication formula)**.** *Let*

$$\Lambda_T = \frac{P^d(T,T_p)A^{a,b}(0)}{A^{a,b}(T)P^d(0,T_p)}$$

be the density in (5.2). Define the formal conditional annuity mapping

$$\alpha_T(s) = \mathbb{E}^A[\Lambda_T \mid S^{a,b}(T) = s].$$

Assume this conditional expectation admits a version such that $h(s) = s\alpha_T(s)$ is twice continuously differentiable and satisfies the integrability condition of Theorem 5.2. Under this regularity assumption, the CMS coupon value admits the representation

$$V_0 = \tau P^d(0,T_p)\Big[h(F) + \int_F^\infty h''(K)C(K)\,dK + \int_{-\infty}^F h''(K)P(K)\,dK\Big], \tag{5.7}$$

where $F = S^{a,b}(0)$, and $C(K)$ and $P(K)$ are payer and receiver swaption prices in annuity units. In practical implementations, α_T is replaced by a tractable approximation such as a swap-yield mapping or Hagan's linear terminal swap-rate mapping [29, 38, 56, 61, 16].

Proof. From (5.2) and the tower property,

$$\mathbb{E}^{T_p}[S^{a,b}(T)] = \mathbb{E}^A[S^{a,b}(T)\Lambda_T] = \mathbb{E}^A[S^{a,b}(T)\alpha_T(S^{a,b}(T))].$$

Applying Theorem 5.2 to $h(s) = s\alpha_T(s)$ gives (5.7). □

This formula identifies the mathematical source of the CMS convexity adjustment. The deterministic term is $h(F) = F\alpha_T(F)$, not simply F unless the annuity-density mapping satisfies $\alpha_T(F) = 1$. The remaining terms integrate payer and receiver swaption prices across strikes, weighted by the curvature of the annuity-adjusted payoff h. After an annuity-mapping approximation has been chosen, a model that fits only the ATM volatility can therefore match the first-order level implied by that mapping while still missing the smile contribution to the convexity adjustment.

5.3 Multi-Curve Correction

In a multi-curve setting, the swap rate contains the forwarding-basis term in (3.3). The following proposition records the first-order implication used as a numerical consistency check under the deterministic-basis approximation of Assumption 0.2; see [59, 33, 27] for the broader multi-curve pricing framework.

Proposition 5.4 (First-order multi-curve correction)**.** *Assume that the fixed and floating payment dates coincide. The forwarding-discounting basis is deterministic, enters only additively through the projected floating-leg numerator in Definition 3.1, and has the small tilt representation*

$$L_{i-1}^{f}(t) - L_{i-1}^{d}(t) = s_0 + s_1(T_i - t) + O(\|s\|^2). \tag{5.8}$$

For a CMS coupon paid at T_p, let $\mathrm{CA}^{\mathrm{SC}}$ be the single-curve convexity adjustment and $\mathrm{CA}^{\mathrm{MC}}$ the multi-curve adjustment under the same OIS discount curve and annuity numeraire. Then

$$\mathrm{CA}^{\mathrm{MC}}(T) = \mathrm{CA}^{\mathrm{SC}}(T) + s_0\Delta_0(T) + s_1\Delta_1(T) + O(\|s\|^2), \tag{5.9}$$

where the sensitivities are

$$\Delta_j(T) = \mathbb{E}^A[\Lambda_T D_j(T)] - D_j(0), \qquad j = 0, 1, \tag{5.10}$$

with

$$\Lambda_T = \frac{P^d(T,T_p)}{A^{a,b}(T)}\frac{A^{a,b}(0)}{P^d(0,T_p)}, \tag{5.11}$$

$$D_0(t) = \frac{\sum_{i=a+1}^{b}\delta_i P^d(t,T_i)}{A^{a,b}(t)}, \qquad D_1(t) = \frac{\sum_{i=a+1}^{b}\delta_i(T_i - t)P^d(t,T_i)}{A^{a,b}(t)}. \tag{5.12}$$

Thus Δ_0 and Δ_1 are the explicit annuity-measure sensitivities of the CMS expectation to a parallel and a linear tilt of the forwarding basis.

Proof. Insert the basis expansion into Definition 3.1. The resulting swap rate has the form

$$S^{a,b}(t) = S^{a,b}_{\mathrm{SC}}(t) + s_0 D_0(t) + s_1 D_1(t) + O(\|s\|^2),$$

because the single-curve floating leg contributes $P^d(t,T_a) - P^d(t,T_b)$, while the two basis perturbations contribute the discounted sums defining D_0 and D_1 after division by the annuity. Because the perturbation is deterministic and additive in

the floating-leg numerator, the annuity $A^{a,b}$, payment-date numeraire and density Λ_T are not perturbed. Substituting the expansion into (5.2), using $\mathbb{E}^{T_p}[X] = \mathbb{E}^{A}[\Lambda_T X]$, and differentiating the convexity adjustment $\mathbb{E}^{T_p}[S^{a,b}(T)] - S^{a,b}(0)$ with respect to (s_0, s_1) at the origin yields (5.9). The sensitivities are finite under Assumptions 0.3 and 0.4. □

The empirical EUR basis in the 2026-03-17 market data is small relative to the level of CMS rates. Consequently, the first-order correction is a useful diagnostic for plain CMS coupons, while callable CMS spread products still require simulation of the joint swap-rate dynamics.

5.4 CMS Spread Pricing

For a CMS spread payoff

$$\Phi\Big(\alpha_1 S^{(1)}(T) - \alpha_2 S^{(2)}(T) + s\Big),$$

one-dimensional static replication is insufficient because the value depends on the joint law of $(S^{(1)}(T), S^{(2)}(T))$, not only on the two marginal smiles. This is the same modelling issue that motivates dynamic market models for families of forward and swap rates [10, 52, 40, 11, 50, 31].

Proposition 5.5 (Instantaneous CMS spread variance)**.** *Let two swap rates have diffusion vectors $\nu_1(t)$ and $\nu_2(t)$ under a common pricing measure, with Brownian instantaneous correlation matrix Γ. For the spread*

$$Y_t = \alpha_1 S_t^{(1)} - \alpha_2 S_t^{(2)} + s,$$

the instantaneous quadratic variation is

$$\frac{d\langle Y\rangle_t}{dt} = \alpha_1^2 \nu_1(t)^\top \Gamma \nu_1(t) + \alpha_2^2 \nu_2(t)^\top \Gamma \nu_2(t) - 2\alpha_1\alpha_2 \nu_1(t)^\top \Gamma \nu_2(t). \tag{5.13}$$

In a one-factor model, $\nu_i(t) = \beta_i(t)\gamma(t)$. Hence

$$\frac{d\langle Y\rangle_t}{dt} = \Big(\alpha_1\beta_1(t) - \alpha_2\beta_2(t)\Big)^2 \gamma(t)^\top \Gamma \gamma(t),$$

so the spread volatility collapses whenever the two swap-rate exposures are nearly proportional after gearing.

Proof. The diffusion of Y is $(\alpha_1\nu_1 - \alpha_2\nu_2)^\top dW$. Taking its quadratic variation gives (5.13). The one-factor statement follows by substituting $\nu_i = \beta_i\gamma$. If the driving Brownian vector has first been whitened so that $\Gamma = I$, the final factor reduces to $\|\gamma(t)\|^2$. □

Equation (5.13) explains why CMS spread products are more discriminating model tests than plain CMS coupons. A single CMS coupon can be largely controlled by its marginal swaption smile through Proposition 5.3. A CMS spread option also requires the covariance term $\nu_1^\top\Gamma\nu_2$. In one-factor Gaussian models the two swap-rate volatility vectors are collinear by construction. Such models can still generate spread volatility when the two loadings differ, but they cannot generate independent cross-tenor movements; when the geared exposures are nearly proportional, the residual spread variance can therefore be very small. Multi-factor QG-SV models relax this restriction by assigning different benchmark tenors and factor loadings to different parts of the term structure.

5.5 Callable CMS Pricing

Proposition 5.6 (Snell envelope for callable CMS products)**.** *Let X_t be the Markov state of the pricing model. Let $0 = T_0 < T_1 < \cdots < T_E$ be the valuation date and issuer exercise dates, and let T_{E+1} denote the final contractual horizon. For $e = 1, \ldots, E$, let $H_e(X_{T_e})$ be the value to the investor if the issuer stops the contract at T_e, including any cashflow paid on stopping. For $e = 0, \ldots, E$, let I_{e+1} denote the $\mathcal{F}_{T_{e+1}}$-measurable value, discounted back to T_e, of contractual cashflows paid after T_e and up to T_{e+1} if the contract has not been cancelled. With $V_{E+1} = 0$, the issuer-callable value satisfies the Snell-envelope backward recursion [57, 66]*

$$V_e(x) = \min\left(H_e(x), C_e(x)\right), \tag{5.14}$$

$$C_e(x) = \mathbb{E}^{\mathbb{Q}}\Big[I_{e+1} + D(T_e, T_{e+1})V_{e+1}(X_{T_{e+1}}) \mid X_{T_e} = x\Big], \tag{5.15}$$

for $e = E, \ldots, 1$. If no cashflow is paid between two adjacent exercise dates, $I_{e+1} = 0$. The time-zero value is

$$V_0 = \mathbb{E}^{\mathbb{Q}}[I_1 + D(0, T_1)V_1(X_{T_1}) \mid X_0],$$

with the convention $I_1 = 0$ if no cashflow occurs before the first exercise date.

Proof. The issuer cancels when the cancellation payment H_e is lower than the continuation cost of the contract. From the investor's perspective, the received

value at an exercise date is therefore the minimum of cancellation and continuation values. If the contract continues from T_e to T_{e+1}, the investor receives the interim cashflows represented by I_{e+1} and then the value-to-go V_{e+1}. At T_E the same formula compares stopping with the value of the final tail I_{E+1}, because $V_{E+1} = 0$. Dynamic programming for the finite set of stopping dates gives the Snell-envelope recursion (5.15). The Markov property of X_t reduces the continuation value to a deterministic function $C_e(X_{T_e})$ of the current state. Discounting the pre-first-exercise cashflows and V_1 gives the stated time-zero value. □

Proposition 5.7 (Issuer-option decomposition)**.** *Let V_0^{NC} be the value of the corresponding non-callable CMS cashflow stream. Suppose that cancellation at a Bermudan stopping time $\tau \in \mathcal{T}$ removes all later cashflows and involves no additional cancellation payment. Assume that $\mathcal{T}$ includes the no-call choice, represented by a terminal cemetery time with $R_\tau = 0$. Let*

$$R_\tau = \sum_{i:T_i>\tau} D(0,T_i)C_i$$

be the discounted value of the cashflows removed if the issuer cancels at stopping time τ. Then the investor value of the issuer-callable product is

$$V_0^{\mathrm{call}} = V_0^{\mathrm{NC}} - \sup_{\tau\in\mathcal{T}} \mathbb{E}^{\mathbb{Q}}[R_\tau]. \tag{5.16}$$

Thus the cancellation right is the Snell-envelope value [57] of a Bermudan option on the remaining CMS tail; the state dependence enters through the admissible stopping time τ.

Proof. For a fixed stopping time τ, the investor receives only cashflows with $T_i \le \tau$. Therefore

$$\mathbb{E}^{\mathbb{Q}}\left[\sum_{i:T_i\le\tau} D(0,T_i)C_i\right] = V_0^{\mathrm{NC}} - \mathbb{E}^{\mathbb{Q}}[R_\tau].$$

Taking the infimum over admissible issuer stopping times gives

$$\inf_\tau\{V_0^{\mathrm{NC}} - \mathbb{E}^{\mathbb{Q}}[R_\tau]\} = V_0^{\mathrm{NC}} - \sup_\tau \mathbb{E}^{\mathbb{Q}}[R_\tau].$$

□

The Longstaff–Schwartz algorithm used in Chapter 6 approximates the continuation function C_e in (5.15) by a least-squares projection on polynomial basis functions of the simulated state [46]. On each simulated path, the regression target is the

discounted value-to-go from the next exercise date (or the later realised cashflow after previous backward induction), not the undiscounted next-step payoff. LSM therefore replaces the exact conditional expectation by a finite sample projection onto a chosen basis. The resulting price is the value of the exercise policy generated by the regression, not the exact Snell envelope unless the continuation function lies in the chosen span and the simulation error vanishes. The numerical price is conditional on this basis choice and has three error sources: time-discretisation error in the simulated QG-SV state, statistical error from the finite path set, and projection error from the chosen basis. The empirical comparison keeps the regression configuration fixed across model families so that changes in basis design are not mistaken for model risk.

5.6 Summary

Plain CMS pricing is a change-of-numeraire and static-replication problem. CMS spread pricing is a joint-distribution problem. Callable CMS pricing is a finite-horizon optimal-stopping problem. The QG-SV model is introduced because it addresses all three issues simultaneously: it fits the swaption smile, gives multi-factor spread dynamics, and remains Markovian enough for regression-based Monte Carlo.

6 Numerical Implementation and Calibration

This chapter documents the numerical choices on which the empirical study of Chapter 7 depends: the calibration objective, the simulation scheme for the multi-factor QG-SV state and CIR variance processes, the Longstaff–Schwartz configuration used for callable pricing, and the variance reduction assumptions. Standard background on Monte Carlo simulation, finite-difference methods, variance reduction and simulation-based exercise for interest-rate models is covered in Glasserman [24], Andersen and Piterbarg [5], and the American-option simulation literature [46, 65, 3]; the chapter therefore records only the methodological choices needed to interpret the thesis results.

The empirical implementation is built in the Open Source Risk Engine (ORE) [55], using QuantLib [68] and QuantExt for market construction, instrument representation and numerical pricing infrastructure. For this thesis, the ORE/QuantExt codebase was extended with the multi-factor QG-SV state process, calibration routines and Monte Carlo pricing components used in the empirical study.

6.1 Calibration Framework

The swaption smiles used for calibration are interpreted in the market normal (Bachelier) volatility convention rather than as lognormal Black volatilities [11, 62, 30]. A normal volatility has units of rate per square-root year, so a one-basis-point volatility error means 10^{-4} in absolute rate-volatility units. Strike labels such as ± 50 bp are also absolute rate spreads around the at-the-money forward, i.e. $K = F \pm 0.0050$. Consequently the RMSE values reported in this thesis are not differences between Black-volatility percentages; each model swaption price is first converted back to the same normal-volatility quote convention and strike grid as the market data. This convention is important in low-rate EUR markets, where a direct lognormal-volatility comparison can be misleading or unavailable for negative strikes.

Calibration is formulated as a weighted least-squares problem in implied volatility space. Let $\mathcal{I}$ index the calibration swaptions with market implied volatilities σ_i^{mkt} and model implied volatilities $\sigma_i^{\mathrm{model}}(\theta)$ obtained from the Monte Carlo price under parameters θ. The bootstrap step at expiry T_n minimises

$$\theta^{(n),*} = \arg\min_{\theta^{(n)} \in \Theta_n} \sum_{i \in \mathcal{I}_n} w_i \left[\sigma_i^{\mathrm{model}}(\theta^{(1)*}, \ldots, \theta^{(n-1)*}, \theta^{(n)}) - \sigma_i^{\mathrm{mkt}} \right]^2, \tag{6.1}$$

with $\theta^{(n)} = (\lambda_k^{(n)}, b_k^{(n)}, \eta^{(n)})_{k=1,\ldots,d}$ and previously calibrated buckets frozen at their

optima. The weights are $w_{\text{ATM}} = 3$ and $w_{\text{off-ATM}} = 1$, which reflects the fact that the callable products in Chapter 7 are most sensitive to the ATM level of the smile, while still requiring a reasonable wing fit. Box constraints are imposed via $\lambda_k^{(n)} \in [10^{-3}, 1]$, a logit transform for $b_k^{(n)}$ with effective range approximately $[0.0067, 0.9933]$, and $\eta^{(n)} \in [e^{-1}, e^3]$ (log-reparametrised). The optimiser is Nelder–Mead simplex [53, 54] with deterministic heuristic starting points: an ATM-volatility start, a lower-volatility/lower-displacement start, the previous expiry's optimum when available, and a wider long-end start for long expiries. A final refinement from the best candidate is attempted when the early-exit threshold has not already been met.

The bootstrap structure follows the calibration algorithm of Chapter 4, Algorithm 1: in the continuous-time piecewise model with fixed structural loadings, the swaption price expiring at T_n depends on $\{\theta^{(1)}, \ldots, \theta^{(n)}\}$ only. The implemented calibration uses the corresponding finite-path, time-discretised Monte Carlo estimator, so the forward bootstrap should be read as a controlled numerical approximation to that local structure rather than as an exact global optimiser. For the reported smile-surface dimension sweeps we additionally use a semi-analytical joint-polish step: the Andersen–Piterbarg QG-SV swaption approximation [5] supplies fast residual evaluations, and a Levenberg–Marquardt least-squares solve [44, 49, 54] moves all expiry buckets together after the sequential bootstrap warm start.

The joint step is regularised with Tikhonov rows that penalise deviations of the factor-volatility loadings from the current LM start [69]. This is not a claim that local optimisation can always find the global minimum in the high-dimensional problem. If lower-dimensional calibrations are embedded in a higher-dimensional configuration, they are treated only as additional LM starting points; the reported fit is the post-optimisation market RMSE of the calibrated candidate with the lowest objective value. Thus the methodology keeps the theoretical nesting argument as guidance for constructing starts, while avoiding a silent substitution that would overstate the observed high-dimensional calibration performance.

6.1.1 Monte Carlo Bootstrap Calibration

The Monte Carlo bootstrap calibrates directly against the same state process, time discretisation and piecewise parameter representation that the callable CMS pricing methods later reuse. This deliberately keeps the calibration objective and the pricing objective on the same numerical surface. Section 7.4 and Proposition A.6 give sufficient conditions under which the semi-analytical frozen-sensitivity approximation

changes the effective variance target; parameters calibrated under that approximation are therefore not treated as automatically valid full Monte Carlo inputs.

Table 6.1: Implemented $d = 2$ MC bootstrap configuration used for the Chapter 7 nested-QG comparison.

Item	**Reported setting**
Calibration expiries	1Y, 2Y, 3Y, 5Y, 7Y, 10Y
Swap tenor	10Y
Strike spreads	-200, -100, -50, 0, $+50$, $+100$, $+200$ bp
Benchmark tenors	$(3.0, 1.5)$ years
Mean reversions	$\kappa = (0.15, 0.45)$
Correlation decay	$\alpha = 0.05$
Variance parameters	$\theta_z = 0.20$, $z_0 = \bar{z} = 1.0$
Calibration paths	1 000 Sobol draws for $T_n \leq 5$Y; 1 500 for longer buckets, with antithetic pairing
Pricing paths	2 000 Sobol Brownian-bridge paths in the frozen valuation runs
Time discretisation	12 steps/year for short buckets and pricing; 16 steps/year for longer calibration buckets
Objective weights	$w_{\text{ATM}} = 3$, $w_{\text{off-ATM}} = 1$
Optimiser	Nelder–Mead simplex with deterministic heuristic starts and early-exit thresholds
Saved diagnostics	per-expiry RMS, ATM forwards, strike residuals, piecewise λ, b, η

The calibrator proceeds expiry by expiry. At bucket n, parameters from earlier intervals are frozen and only the current interval values of $\lambda_k^{(n)}$, $b_k^{(n)}$ and $\eta^{(n)}$ are re-optimised. The result is a piecewise parameter term structure that can be used immediately for pricing. The calibration diagnostics contain the per-expiry RMS strip, ATM forwards, the full strike-level smile residuals, and flattened parameter vectors; Appendix B tabulates the saved values.

For the progressive scenarios, the same calibrator is reused for the three nested model families of Chapter 7. In the Gaussian scenario, the displacement parameters are fixed to zero and η is held at a near-zero floor; in the skew-only scenario, η is held at the floor while the displacement terms remain free; in the full QG-SV scenario, all $2d + 1$ parameters per interval are optimised. This avoids changing the numerical method between Gaussian, skew and stochastic-volatility runs, so the comparison isolates added model degrees of freedom rather than numerical-method differences.

6.2 Monte Carlo Simulation of the QG-SV State

The QG-SV state vector $(x(t), y(t), z(t))$ is simulated on the time grid specified by the calibration or pricing run. The reported valuation runs use 12 steps per year for pricing and for short calibration buckets; the adaptive bootstrap increases this to 16 steps per year for longer calibration buckets. The key numerical choice is the discretisation of the CIR variance process

$$dz(t) = \theta_z(\bar{z} - z(t))\,dt + \eta\sqrt{z(t)}\,dW^z(t), \tag{6.2}$$

which is non-Lipschitz at $z = 0$ and can become negative under naive Euler. The implemented process does not use a full-truncation Euler drift step. Instead it applies the Andersen–Piterbarg moment-matched Gaussian approximation to the CIR transition [5, 24, 47]: for a step of length Δt,

$$\begin{aligned} m_i &= \bar{z} + (z_i^+ - \bar{z})e^{-\theta_z \Delta t}, \\ v_i &= z_i^+ \frac{\eta^2}{\theta_z} e^{-\theta_z \Delta t}\left(1 - e^{-\theta_z \Delta t}\right) + \bar{z}\frac{\eta^2}{2\theta_z}\left(1 - e^{-\theta_z \Delta t}\right)^2, \\ z_{i+1} &= \max\{m_i + \sqrt{\max(v_i, 0)}\, Z_i^z, 0\}, \qquad z_i^+ = \max(z_i, 0). \end{aligned} \tag{6.3}$$

The state factors x_k and auxiliary variances y_{kl} are propagated by Euler steps using the displaced QG-SV drift and diffusion derived in Chapter 4. Numerical safeguards bound x_k to $[-1, 1]$ and y_{kl} to $[-0.01, 0.05]$, and non-finite state updates retain the previous state. These safeguards are not part of the continuous-time model; they define the finite-grid approximation used in the empirical study and introduce a possible discretisation bias if paths reach the bounds. The reported diagnostics indicate that the bounds are far outside the typical calibrated state range, so they are treated as reproducibility safeguards rather than active modelling assumptions. We adopt the recommendation of [5] to take $\rho = 0$ between z and the rate factors, which removes the cross-correlation in the discretisation; Chapter 7 reports the resulting smile fit on the EUR data set.

The random inputs are generated with a Sobol Brownian-bridge construction and, in the bootstrap calibration, antithetic pairing [67, 24]. The callable pricing runs reuse the Sobol Brownian bridge but do *not* add the antithetic mirror path, because Bermudan exercise breaks the symmetry that antithetic pairing exploits and would require additional regression bookkeeping. The exact path-count settings are reported in Table 6.1, while their implications for Monte Carlo uncertainty are discussed in Chapter 7.

6.3 Longstaff–Schwartz for Callable Pricing

Callable CMS, callable CMS-spread and Bermudan swaption prices in Chapter 7 are computed by Longstaff–Schwartz regression [46] on the simulated paths, with the duality literature providing the standard benchmark for understanding the possible bias of regression-based exercise policies [65, 3]. At each exercise date T_e the continuation value is regressed on an additive univariate polynomial basis of the selected rate-factor state variables

$$C(T_e, x) \approx \sum_{r=1}^{R} \sum_{k=0}^{p} \alpha_{r,k}\, x_r^k, \tag{6.4}$$

with $R = \min(d, 3)$ regressors (the rate factors only; the variance factor z is not used as a regressor because it is approximately uncorrelated with the rate innovations under $\rho = 0$) and polynomial order $p = 4$. The regression is performed on in-the-money paths only and is solved via singular-value decomposition of the design matrix, which is robust against near-collinear regressors at long expiries [26]. Omitting z is a basis restriction rather than a theorem: even with $\rho = 0$, the variance factor can affect future option values through conditional variance. Chapter 7 treats this as part of the Monte Carlo uncertainty and LSM-basis limitation. The regression estimates the continuation value; the exercise comparison then follows the contract's optionality direction. For a Bermudan swaption held by the option owner, the realised path value uses the maximum of exercise and continuation values. For an issuer-callable CMS product valued from the investor perspective, the issuer chooses the lower investor value, so the recursion uses the minimum of the call-settlement value and the continuation value; equivalently, the issuer's cancellation option can be valued with a maximum on the issuer's saving. The same regression configuration is used for the LGM, G2++ and QG-SV variants in Table 7.3 where the numerical methods allow it; see [25] for the dependence of LSM accuracy on basis dimension and path count.

6.4 Calibration Diagnostics

The calibration diagnostics record per-expiry RMS errors, ATM forward swap rates, strike-level smile residuals, and the flattened piecewise parameters $(\lambda_k^{(n)}, b_k^{(n)}, \eta^{(n)})$. Appendix B reproduces these quantities for the $d = 2$ calibration that drives the empirical results in Chapter 7. The optimisation diagnostics also record Nelder–Mead progress, which helps identify local minima and verify that the reported full QG-SV

calibration fits all expiry buckets to within 3.1 bp.

6.5 Summary

The numerical study reduces to four concrete choices that drive every empirical number in Chapter 7: a Nelder–Mead bootstrap with ATM-weighted least squares in implied volatility space, an optional Andersen–Piterbarg semi-analytical joint LM–Tikhonov polish for surface sweeps, a moment-matched CIR variance step on Sobol Brownian-bridge grids, and a Longstaff–Schwartz regression with polynomial basis in the rate factors. The MC bootstrap uses antithetic Sobol pairing; the callable pricing runs do not. The next chapter applies this methodology to the empirical comparison.

7 Comparative Analysis

This chapter compares CMS instrument prices across the model hierarchy developed in Chapter 4. The aim is not a universal model ranking, but a controlled case study of which added model features remain economically visible after aligning market inputs, calibration targets and downstream pricing tasks as far as the available methods allow. The comparison is judged by calibration RMSE, movement of a common European anchor, callable NPV or exercise premium, and the degree to which each table is confounded by benchmark choice, recalibration or factor count.

The chapter proceeds from calibration quality to three empirical comparisons: Table 7.3 gives a cross-family fixed-strike benchmark, Table 7.9 gives the most controlled nested Gaussian/QG+Skew/QG-SV comparison, and Table 7.11 uses PCA to interpret factor sensitivity. The remaining sections provide mechanism checks, the product matrix, the semi-analytical versus Monte Carlo calibration issue, computational cost and limitations.

Table 7.1 fixes the economic trade specifications used in the pricing comparisons. All trades are EUR 10M notional, collateralised and discounted on the OIS curve, with EUR-EURIBOR-6M used for floating-rate projection where applicable.

Table 7.1: Priced trade specifications used in the empirical comparison.

Group	**Trades**	**Economic specification**
Vanilla swaption anchors	European payer swaptions 1Y×10Y, 5Y×10Y, 10Y×10Y; Bermudan payer swaptions 1Y×5Y, 2Y×10Y, 5Y×10Y, 10Y×10Y, 5Y×20Y	Physical payer swaptions on annual-fixed/6M-floating EUR swaps; fixed rate 2.7%, 30/360 fixed leg and Act/360 floating leg. Bermudan exercise is annual after the non-call period.
Analytical CMS controls	CMS 10Y swap, CMS 10Y–2Y spread swap, capped CMS 10Y swap	Receive CMS 10Y versus fixed 2.7%; receive $\max(\text{CMS}_{10Y} - \text{CMS}_{2Y} + 10\text{ bp}, 0)$ versus fixed 0%; receive $\min(\text{CMS}_{10Y}, 4\%)$ versus fixed 2.5%.
Standard callable CMS portfolio	5nc10Y callable CMS 10Y swap and 5nc10Y callable CMS 10Y–2Y spread	Issuer-callable annual Bermudan structures. The single-index trade receives CMS 10Y and pays fixed 2.7%; the spread trade receives $\max(\text{CMS}_{10Y} - \text{CMS}_{2Y} + 10\text{ bp}, 0)$ and pays fixed 0%.
Stress decomposition trade	1nc10Y callable CMS 10Y–2Y spread	Issuer-callable annual Bermudan structure used only for the progressive decomposition: receive quarterly $3\max(\text{CMS}_{10Y} - \text{CMS}_{2Y}, 0)$ and pay semi-annual EUR-EURIBOR-6M minus 100 bp.

7.1 Calibration Analysis

The first diagnostic is the ability of each model variant to reproduce the EUR swaption implied volatility surface observed on 17 March 2026. All calibrations target the same strip of European swaptions with expiries 1Y, 2Y, 3Y, 5Y, 7Y, and 10Y into the 10Y swap tenor and strikes spanning ATM±200 bp, using the MC bootstrap procedure of Algorithm 1. The valuation runs use 1 000 Sobol draws for short buckets and 1 500 for longer buckets, with antithetic pairing, and the ATM weight $w_{\text{ATM}} = 3$ versus $w_{\text{off-ATM}} = 1$ in the cost function.

All reported results use EUR curve and swaption-volatility quote levels from the 2026-03-17 market snapshot and the same 2026-03-17 date as the valuation date. The snapshot contains EUR ESTER/OIS discount quotes, Euribor forwarding quotes and EUR normal swaption ATM/smile quotes.

The wider EUR ATM normal-volatility surface used in the empirical study is shown in Figure 7.1. In the market configuration, the EUR swaption volatility key is attached to the EUR CMS swap indices, whose 2Y and longer conventions reference the EUR-EURIBOR-6M floating leg. The figure therefore records the ATM slice of the same EUR swaption volatility input that anchors CMS replication prices and the QG-SV swaption calibration. The QG-SV calibration itself uses the full smile around this ATM backbone, not only the ATM quotes.

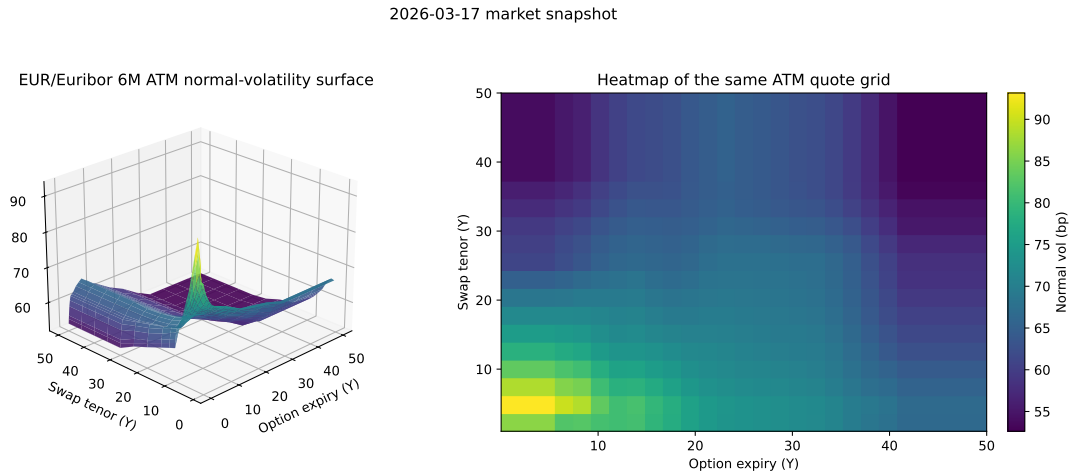

Figure 7.1: EUR/Euribor 6M ATM swaption normal-volatility surface from the market snapshot dated 2026-03-17. The left panel shows the surface view and the right panel shows the same data as a heatmap. Values are normal volatilities in basis points, extracted from the EUR ATM swaption quotes in the market input.

7.1.1 Smile Fit Quality

Table 7.2 summarizes the root-mean-square (RMS) implied volatility error across expiries for the three model variants in the two-factor ($d = 2$) specification.

Table 7.2: Calibration RMS errors (bp of implied volatility) by swaption expiry for the three model variants. Two-factor specification ($d = 2$), MC bootstrap with adaptive 1 000–1 500 Sobol draws and antithetic pairing.

Expiry	MF Gaussian	QG+Skew	Full QG-SV
1Y	10.6	6.1	0.5
2Y	8.7	3.4	1.6
3Y	7.3	3.5	1.8
5Y	4.2	2.1	3.1
7Y	3.7	2.2	2.4
10Y	2.9	1.6	1.7

All three models achieve ATM fit errors below 1 bp, reflecting the elevated ATM weight (3×) in the calibration cost function. Accurate ATM fit is important for consistent CMS pricing because the convexity adjustment is anchored to at-the-money swaption volatilities.

The Gaussian model exhibits the largest residuals at the short end, with 10.6 bp at 1Y and 8.7 bp at 2Y, before improving to below 5 bp for longer expiries. Without displacement or stochastic volatility parameters, the optimizer lacks sufficient degrees of freedom to simultaneously match ATM and wing volatilities, although the residuals remain moderate on the 2026-03-17 surface.

The QG+Skew and Full QG-SV models reduce these errors substantially. Among the reported local-optimisation candidates, the full QG-SV specification gives the lowest overall RMSE (2.00 bp) and keeps every expiry-RMS bucket within 3.1 bp, with the 5Y bucket as the largest residual. The displacement parameter $b_k(t)$ accounts for a large part of the improvement relative to the Gaussian restriction, while stochastic volatility sharpens the fit across the whole strip rather than merely correcting a single outlier. Figure 7.4 overlays the three model smiles against market data, and Figure 7.5 summarizes the calibration quality across models.

7.1.2 Strike-Level Diagnostics

The calibration error is not uniformly distributed across strikes. ATM strikes are fitted most accurately due to elevated weighting, while most wing errors stay in the low single digits. In the full QG-SV run the largest residuals occur in the far

high-strike wings of the 5Y, 7Y, and 10Y smiles, with the worst point around 6.2 bp. Out-of-the-money puts and calls are both captured materially better than in the Gaussian restriction, indicating that the displaced-diffusion and stochastic-volatility parameters improve the smile shape rather than merely the ATM backbone.

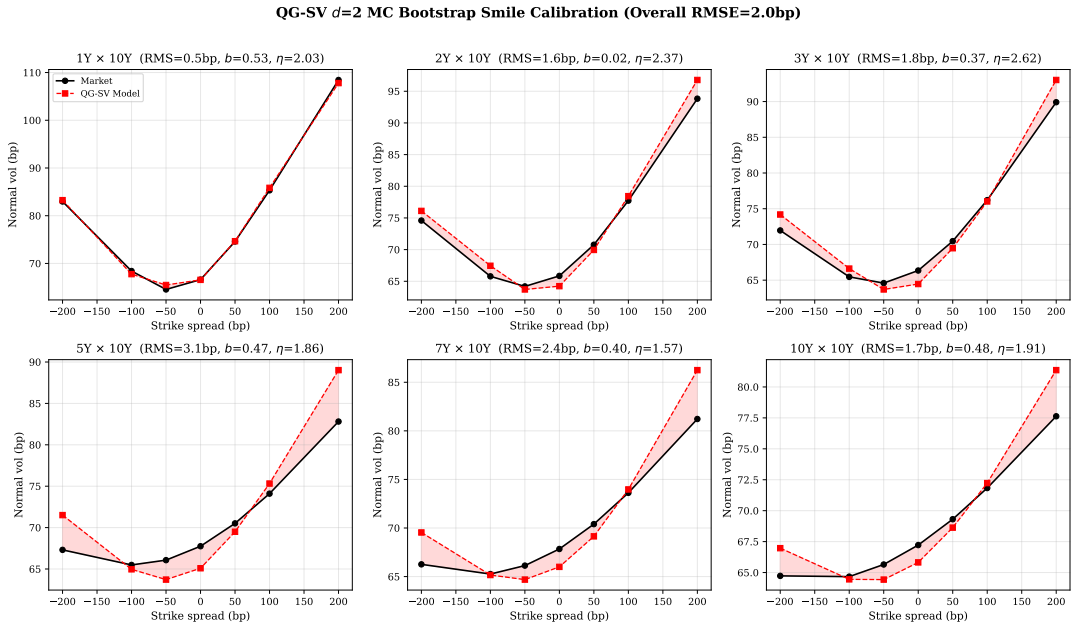

Figure 7.2: MC bootstrap smile fit across swaption expiries. Market implied volatilities (circles) versus model-calibrated volatilities (lines) for strikes spanning ATM±200 bp. The full QG-SV model ($d = 2$) fits the 2026-03-17 smile to 2.00 bp aggregate RMSE, with all expiry buckets below 3.1 bp.

Remark 7.1 (Calibration Stability)**.** The sequential bootstrap (Algorithm 1) calibrates parameters expiry by expiry, freezing earlier intervals. This design keeps the Gaussian parameters $\lambda_k(t)$ fixed during subsequent smile calibration, providing a stable ATM backbone. The risk is that a poor fit at one expiry can propagate forward. Multi-start initialization mitigates but does not eliminate this risk; in the reported valuation run, the full QG-SV strip avoids any dominant single-expiry local minimum, but the result should still be read as the best reported local optimum rather than proof of a global optimum.

7.2 Fixed-Strike Swaption Benchmark Across Model Families

Before turning to CMS products, it is useful to anchor the cross-family comparison on plain swaptions only. Table 7.3 collects the common 5Y×10Y European anchor

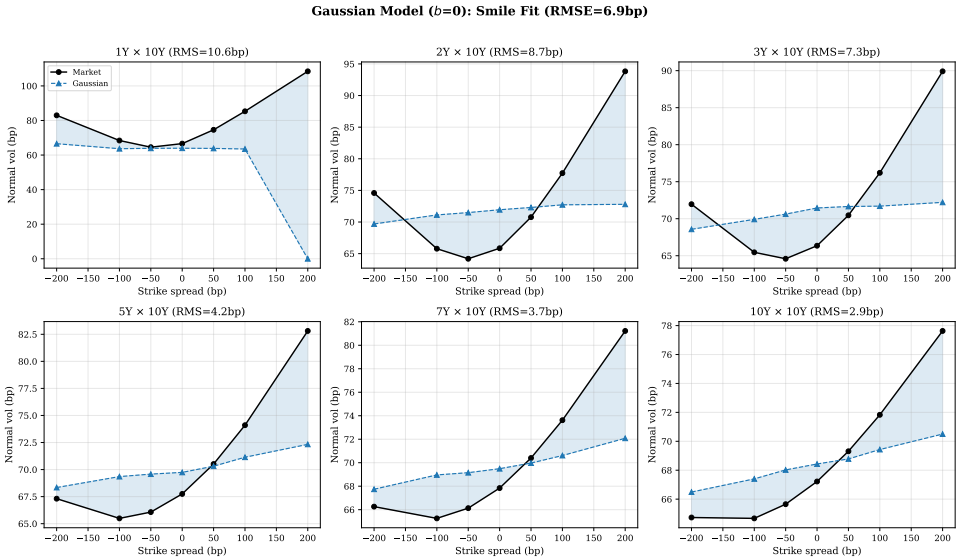

Figure 7.3: Gaussian-restricted calibration overlay. The flat-smile restriction keeps the ATM backbone close but cannot reproduce the observed strike curvature, motivating the skew and stochastic-volatility extensions in Figure 7.4.

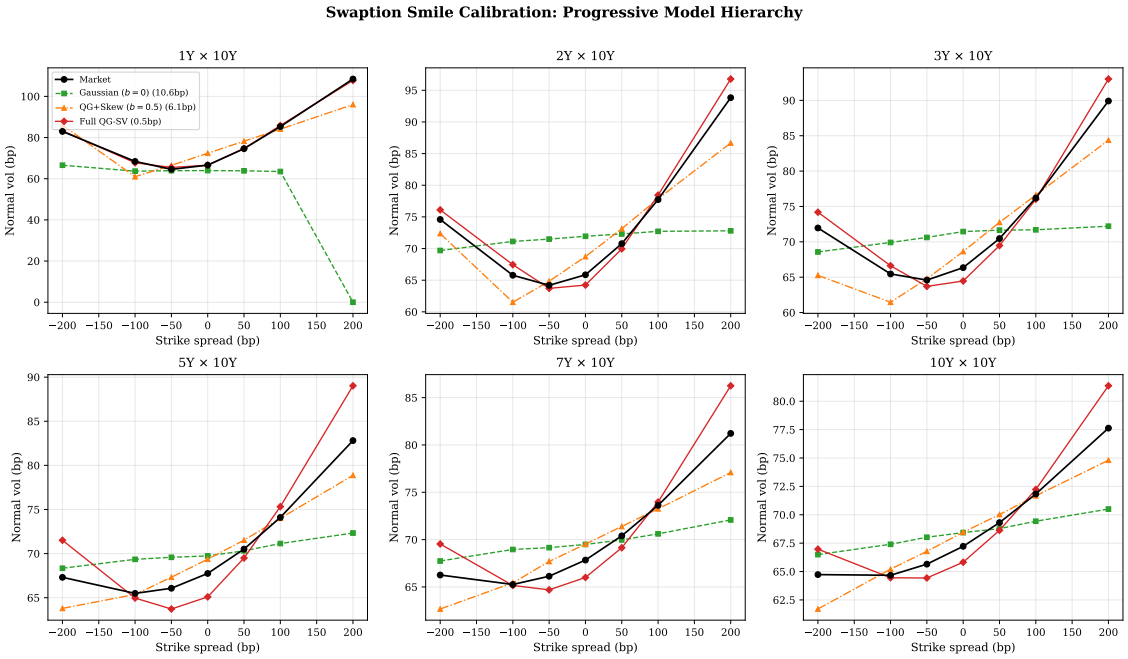

Figure 7.4: Swaption smile calibration across the progressive model hierarchy. Market implied normal volatilities (black circles) are compared with calibrated model smiles for Gaussian ($b = 0$, blue squares), QG+Skew ($b = 0.5$, orange triangles), and full QG-SV (green diamonds) at six expiries into the 10Y swap tenor. Per-expiry RMS errors are shown in parentheses. The Gaussian model produces a flat smile, QG+Skew captures the dominant skew, and full QG-SV gives the lowest aggregate residual among the reported calibration candidates.

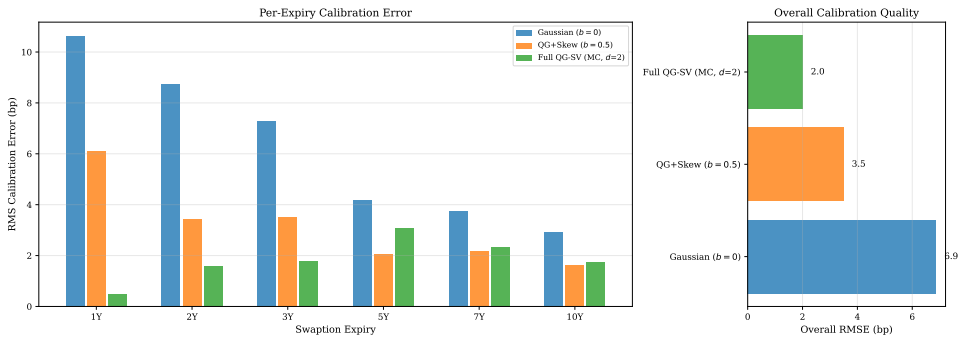

Figure 7.5: Left: per-expiry RMS calibration errors (bp) across the nested $d = 2$ MC model variants. Right: overall RMSE comparison for the same three-variant calibration set. The full QG-SV MC model gives the lowest reported RMSE (2.00 bp), improving on QG+Skew (3.50 bp) and the Gaussian restriction (6.87 bp).

from the Bermudan-swaption runs, the Bermudan basket, and the matching 5Y×10Y exercise premium across LGM, G2++, and MC-calibrated QG-SV. All rows use the fixed 2.7% strike. The G2++ row remains a diagnostic benchmark because it is calibrated and priced under a separate constant-parameter finite-difference setup, with calibration targets selected from the same strike and benchmark basket as the reported swaptions.

Table 7.3: Fixed-strike swaption benchmark across model families. The 5Y×10Y European row gives the common vanilla anchor, while the 5Y×10Y exercise premium is the primary robust comparison metric. Prices in EUR, 10M notional.

Row	**Structure**	**LGM**	**G2++**	**QG-SV**
European	5Y×10Y NPV	724 984	702 341	719 021
Bermudan	1Y×5Y NPV	220 316	164 743	240 802
Bermudan	2Y×10Y NPV	690 853	597 392	811 744
Bermudan	5Y×10Y NPV	838 655	774 584	940 619
Bermudan	10Y×10Y NPV	862 762	826 433	846 553
Bermudan	5Y×20Y NPV	1 482 499	1 397 383	1 436 539
Benchmark	5Y×10Y premium	15.7%	10.3%	30.8%

The 5Y×10Y European anchor is close across the three model columns, while the corresponding Bermudan exercise premium is not ordered by nominal model complexity. In this calibration, G2++ gives a 10.3% 5Y×10Y premium, below LGM's 15.7%, whereas QG-SV gives 30.8%. The table therefore supports the narrower conclusion that Bermudan exercise value is highly calibration- and dynamics-dependent; it does not by itself rank the model families.

7.2.1 G2++ Two-Factor Gaussian Comparison

To isolate the effect of multi-factor dynamics from stochastic volatility, we calibrate the standard G2++ two-factor Gaussian short-rate model [11] to the same EUR market data and price the Bermudan swaption structures from Table 7.3. The G2++ model

$$r(t) = \varphi(t) + x(t) + y(t), \qquad \begin{cases} dx = -a\,x\,dt + \sigma\,dW^1, \\ dy = -b\,y\,dt + \eta_G\,dW^2, \\ \langle dW^1, dW^2 \rangle = \rho\,dt \end{cases} \tag{7.1}$$

provides a two-factor Gaussian specification with five constant parameters: mean reversions a and b, volatilities σ and η_G, and correlation ρ. The deterministic shift $\varphi(t)$ is chosen to match the initial term structure exactly. The notation η_G is used

here to distinguish the second Gaussian factor volatility from the QG-SV variance vol-of-vol η.

The G2++ model uses the same multi-curve market data convention as the other benchmark rows: the OIS (ESTER) curve serves as the discount curve, while a separately bootstrapped EUR-EURIBOR-6M curve provides the forecast rates. The OIS–Euribor basis is approximately 21 bp at the 10Y tenor. Calibration uses analytical European swaption prices with multi-start Levenberg–Marquardt optimization, while Bermudan pricing uses a two-factor finite-difference method consistent with the separate discount and forecast curves.

The reported G2++ row is calibrated to five fixed-strike European swaption targets: 1Y×5Y, 2Y×10Y, 5Y×10Y, 10Y×10Y and 5Y×20Y at the common 2.7% strike. The target normal volatilities are linearly interpolated from the market smile at each expiry–tenor point. The fitted parameters are $a = 0.843$, $\sigma = 0.0338$, $b = 0.0110$, $\eta_G = 0.00728$ and $\rho = -0.905$. The analytical calibration objective matches these five strike-relevant European targets closely, while the reported NPVs are obtained by finite-difference repricing for consistency with the Bermudan calculations. This target selection should not be read as a global smile fit.

Table 7.3 records the common 5Y×10Y exercise premium alongside the Bermudan NPV basket across LGM, G2++ (FDM pricing), and QG-SV MC. The exercise premium $\pi = V_{\text{Berm}}/V_{\text{Eur}} - 1$ normalizes part of the curve-construction and numerical-method difference between the G2++ and QG-SV setups. For the 5Y×10Y structure, G2++ gives a 10.3% premium, compared with 15.7% for LGM and 30.8% for QG-SV. The result shows that the chosen two-factor Gaussian calibration does not reproduce the larger QG-SV exercise premium. It does not decompose callable CMS spread value into separate Gaussian-factor and stochastic-volatility contributions, because no like-for-like G2++ CMS valuation method is implemented here.

7.3 Progressive Model Complexity

The model nesting hierarchy of Proposition 4.13 enables a controlled experiment: by progressively activating model features while keeping the calibration data and numerical method identical, we isolate the pricing impact of each feature on CMS instruments. We consider three levels:

(i) the MF Gaussian row fixes $b_k \equiv 0$ and takes $\eta \to 0$, giving the flat-smile restriction calibrated to the same swaption surface with ATM-heavy weighting;

(ii) the QG+Skew row takes $\eta \to 0$ while calibrating b_k, so that the same calibration

set is fitted with skew but no stochastic-volatility curvature;

(iii) the Full QG-SV row calibrates all parameters, using the same calibration set with both skew and CIR stochastic volatility.

All three variants use the two-factor ($d = 2$) specification of Definition 4.6 and identical frozen-pricing settings: 2 000 Sobol Brownian-bridge paths and 12 time steps per year. The calibration bootstrap discussed in Section 7.1 uses the adaptive 1 000–1 500 draw grid with explicit antithetic pairing.

7.3.1 European Swaption Prices

As a consistency check, we first examine European swaption prices across the three models. This is not a strict invariance test. The three variants share the same calibration set and the same Monte Carlo pricing method, but the weighted objective, finite calibration residuals, and sequential bootstrap all allow a single instrument to move across variants. The purpose of Table 7.4 is therefore to verify that the models remain anchored to the same vanilla market before we interpret the exotic-price dispersion.

Table 7.4: European swaption prices (EUR) across nested QG model variants. 5Y expiry into 10Y swap, ATM strike, EUR 10M notional. All models are calibrated to the 2026-03-17 smile data.

Model Variant	**European Swaption NPV (EUR)**
MF Gaussian	750 032
QG+Skew	731 264
Full QG-SV	727 866

Table 7.4 shows that the nested variants remain closely anchored: the 5Y×10Y European value moves from 750 032 EUR in the Gaussian restriction to 727 866 EUR in the full QG-SV model, a decline of about 3.0%. The models therefore share a common vanilla backbone, while the callable-product dispersion in the following sections reflects structural model effects and exercise dynamics more than residual vanilla misalignment.

7.3.2 Analytical CMS Control Products

Before examining the model-dependent products, we establish an important baseline: the non-callable CMS control rows produce identical NPVs across model variants

because the same analytical CMS pricing inputs are used in every column. Table 7.5 reports the NPVs for three representative plain CMS products.

Table 7.5: NPVs of plain CMS products (EUR) across model variants. All values are identical to the cent because the analytical CMS coupon pricing inputs are held fixed across model columns.

Product	**NPV (EUR)**	**Pricing Method**
CMS 10Y Swap	655 931	LinearTSR
CMS Spread Swap (10Y–2Y)	941 533	Brigo–Mercurio
Capped CMS 10Y	474 350	LinearTSR

This equality is a consequence of the pricing methodology rather than evidence that dynamic model choice is irrelevant in general. Plain CMS coupons are valued using the *linear terminal swap rate* (LinearTSR) formula of [29]. This is a practical implementation of the one-dimensional static-replication logic in Proposition 5.3: once the annuity mapping has been fixed, the initial curve and market-implied swaption smile determine the CMS convexity adjustment without simulating a separate dynamic interest-rate model. CMS spread coupons are valued using the Brigo–Mercurio analytical spread approximation. It is not an exact one-dimensional replication formula, but a joint-distribution approximation based on initial swap rates, marginal swaption volatilities and a specified spread-correlation input [11]. In both cases the method is deliberately market-surface based: the model variants in this chapter are not used to generate future exercise decisions or pathwise spread dynamics for these non-callable rows. The equality of the reported NPVs is therefore a consistency check on the fixed analytical valuation stack. Dynamic model choice enters only when Bermudan exercise, pathwise spread dynamics or model-implied future distributions are required.

7.3.3 Callable CMS Spread Pricing

We now turn to the central comparative result. The test instrument is the stress-decomposition trade in Table 7.1: a 10-year issuer-callable CMS 10Y–2Y spread swap that receives quarterly $3\max(\mathrm{CMS}_{10Y} - \mathrm{CMS}_{2Y}, 0)$, pays semi-annual EUR-EURIBOR-6M minus 100 bp, and has annual Bermudan call rights from year 1. This product combines the three CMS mechanisms derived in Chapter 5: CMS convexity adjustments depend on the full swaption smile, the long-short spread structure depends on the covariance term in Proposition 5.5, and Bermudan exercise decisions subtract the issuer option in Proposition 5.7. Note that this product specification

differs from the 5nc10Y callable CMS spread with 1× gearing and a 10 bp coupon spread used in the product comparison matrix (Table 7.13); the leveraged structure is chosen here to amplify and isolate model effects in the progressive decomposition.

Table 7.6 presents the pricing results across the three model variants.

Table 7.6: Callable CMS 10Y–2Y spread swap pricing across model variants for the stress decomposition trade: receive quarterly $3\max(\mathrm{CMS}_{10Y} - \mathrm{CMS}_{2Y}, 0)$, pay semi-annual EUR-EURIBOR-6M minus 100 bp, 1nc10Y structure. The underlying NPV is common across columns because the analytical CMS-spread valuation inputs are held fixed.

Model	**Callable NPV (EUR)**	**Option Value (EUR)**	**Option/Underlying**
Underlying (all models)	1 057 181	—	—
MF Gaussian	388 015	669 165	63.3%
QG+Skew	404 353	652 828	61.8%
Full QG-SV	392 780	664 401	62.8%

The underlying swap NPV of 1 057 181 EUR is common to all models as it depends only on the initial yield curve, not on volatility dynamics. Consistent with (5.16), the callable NPV equals the underlying NPV minus the option value granted to the issuer, so a higher cancellation-option value reduces the callable NPV received by the investor.

The results reveal material but moderate model dependence. The Gaussian model produces the lowest callable NPV (388 015 EUR) with the highest option value (669 165 EUR, or 63.3% of underlying). Adding displaced-diffusion skew increases the callable NPV by 16 337 EUR to 404 353 EUR, while the full QG-SV model produces an intermediate callable NPV of 392 780 EUR.

7.3.4 Decomposing Model Effects

We decompose the total model effect into two components: the *skew effect* (Gaussian → QG+Skew) and the *smile effect* (QG+Skew → Full QG-SV). Table 7.7 summarizes the decomposition.

Table 7.7: Decomposition of model effects on callable CMS spread pricing. Each effect is measured as the change in callable NPV relative to the simpler model.

Model Effect	**Δ Callable NPV (EUR)**	**Relative Change**
Skew effect (Gaussian → QG+Skew)	+16 337	+4.2%
Smile effect (QG+Skew → Full QG-SV)	−11 573	−2.9%
Total (Gaussian → Full QG-SV)	+4 764	+1.2%

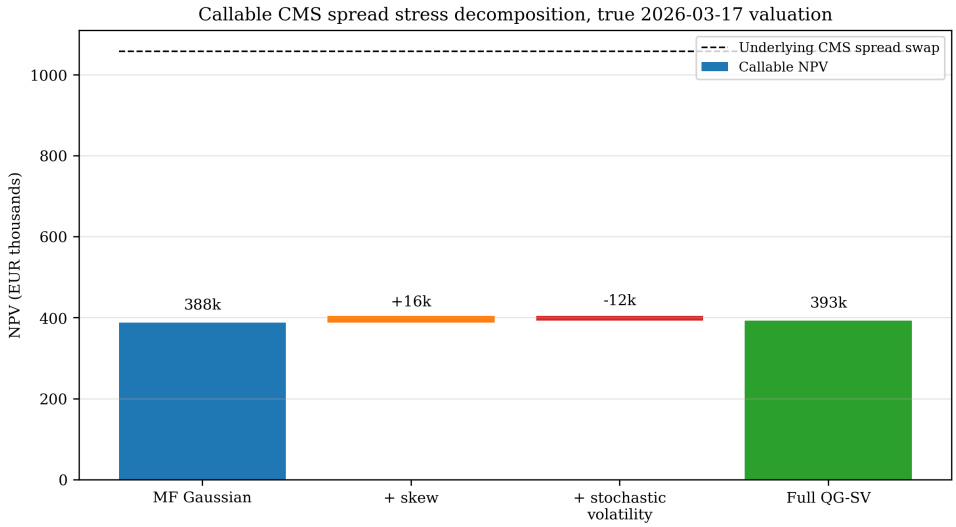

Figure 7.6: Progressive model complexity waterfall: decomposition of the callable CMS spread swap NPV across the three model variants. The underlying NPV (1 057 181 EUR) is common across columns under the fixed analytical CMS-spread valuation inputs. Skew increases the callable NPV by +4.2%, while stochastic volatility adjusts it by −2.9%.

The skew effect (+16 337 EUR, or +4.2% on callable NPV) is the largest single effect in the progressive decomposition. The displacement parameter $b_k(t)$ introduces negative skew into the swap-rate distribution, lowering the value of the issuer's Bermudan cancellation right. CIR stochastic volatility then partially offsets this effect (−11 573 EUR, or −2.9%) by reintroducing tail weight. The net Gaussian-to-QG-SV effect is modest, but the intermediate rows show that skew and kurtosis remain economically active model-risk channels.

Remark 7.2 (Contrast with European Swaptions)**.** The European anchor in Table 7.4 drifts by about 3% across the three variants, while the stress-amplified callable CMS spread NPV in Table 7.6 also moves by a few percent in net but shows larger opposing intermediate effects. The lesson is that callable CMS spread products can amplify skew and kurtosis channels, but the sign and net size of those channels are market- and product-dependent.

Observation 7.3 (Option value is not monotone in model complexity)**.** The option value is not monotone across the model hierarchy: the Gaussian model produces the highest option value (669 165 EUR), QG+Skew the lowest (652 828 EUR), and full QG-SV an intermediate value (664 401 EUR). This is consistent with the competing effects of skew and kurtosis, and it prevents the results from being interpreted as a simple complexity ranking.

7.3.5 Callable CMS Swap Pricing

We now examine the callable CMS swap (without spread) to isolate the model effect on single-index CMS optionality. The test instrument is the standard 5nc10Y callable CMS 10Y swap in Table 7.1: receive annual CMS 10Y coupons and pay fixed 2.7%, with annual Bermudan call rights after the five-year non-call period. Table 7.8 presents the pricing results across model variants and dimensions.

Table 7.8: Callable CMS 10Y swap pricing across model variants and factor dimensions. 5nc10Y structure, receive annual CMS 10Y and pay fixed 2.7%, annual Bermudan call right after year 5.

Model Variant	Callable NPV (EUR)	Implied Option Value (EUR)
LGM	138 373	517 558
QG-SV ($d = 2$)	56 203	599 728
QG-SV ($d = 3$)	45 462	610 469
QG-SV ($d = 4$)	33 850	622 081

The callable CMS swap exhibits material model dependence: the callable NPV ranges from 34 K (QG-SV $d = 4$) to 56 K (QG-SV $d = 2$) within the QG-SV family, while LGM gives a higher value of 138 K. The implied option value (cancellation right) ranges from 600 K to 622 K across QG-SV dimensions and is 518 K under LGM, indicating that the one-factor benchmark assigns a lower issuer cancellation value for the single-index callable CMS swap in this empirical setup.

The dependence on the number of factors is noteworthy: the $d = 4$ specification produces the lowest callable NPV (highest option value), while the $d = 2$ row is highest within the QG-SV family. This reflects the interaction between the number of independent rate drivers and the calibration quality: with more factors, the model has additional degrees of freedom but also more parameters to calibrate from the same market data, and the bootstrap may settle into different local optima.

7.3.6 Nested QG Smile Comparison on the Full Callable CMS Portfolio

The leveraged 1nc10Y callable CMS spread of Table 7.6 is useful for mechanism identification, but it is intentionally stress-amplified. Table 7.9 therefore repeats the nested Gaussian → QG+Skew → QG-SV comparison on the standard 5nc10Y callable CMS portfolio specified in Table 7.1. This is the same portfolio specification used in the full product matrix, but it is distinct from the stress-amplified 1nc10Y decomposition trade. All three rows use the same full-smile calibration family (1Y, 2Y, 3Y, 5Y, 7Y, 10Y into 10Y). To keep the comparison controlled, all three rows are repriced with frozen piecewise parameter term structures recovered from the corresponding calibration runs, while the portfolio and Monte Carlo pricing method are held fixed.

Table 7.9: Nested QG smile comparison on the full callable CMS portfolio. All three variants share the same full-smile calibration family and the same 5nc10Y callable CMS portfolio: receive CMS 10Y versus fixed 2.7% for the single-index trade, and receive $\max(\text{CMS}_{10Y} - \text{CMS}_{2Y} + 10\text{ bp}, 0)$ versus fixed 0% for the spread trade. Prices in EUR, 10M notional.

Metric	**MF Gaussian**	**QG+Skew**	**Full QG-SV**
Calibration RMSE (bp)	6.87	3.50	2.00
5Y×10Y European	750 032	731 264	727 866
Callable CMS Swap	75 189	81 749	56 203
CMS Swap Option Value	580 742	574 182	599 728
Callable CMS Spread	698 226	692 337	695 659
CMS Spread Option Value	243 308	249 196	245 875

Observation 7.4 (Nested smile hierarchy on the full portfolio)**.** The 5Y×10Y European anchor drifts only modestly across the three variants (750 K → 728 K), so the full-portfolio comparison remains attached to a common vanilla backbone. The two callable products then react differently. The single-index callable CMS swap is highest under the skew-only row and lowest under full QG-SV, with option values between 574 K and 600 K. By contrast, the callable CMS spread is very stable, with NPVs clustered in a narrow 692–698 K band and option values between 243 K and 249 K. Hence improved smile fit changes the decomposition of option value, but on this standard 5nc10Y portfolio the standard-product prices are less sensitive than the stress-amplified 1nc10Y trade.

Relative to the stress-amplified 1nc10Y decomposition instrument, the standard 5nc10Y portfolio compresses the nested-QG dispersion substantially. Within the reported local calibrations, smile RMSE improves along the nested sequence and the vanilla anchor remains close, but the callable products do not respond uniformly. Additional smile structure changes callable risk in a product-specific way rather than making the full QG-SV row uniformly preferable for every payoff.

7.3.7 LGM Benchmark for Callable CMS Products

The comparison between the LGM (Gaussian one-factor) model and the QG-SV model is informative for both single-index and spread-contingent callable CMS products. The LGM comparison uses a pathwise Monte Carlo valuation in which CMS rates are evaluated stochastically at exercise dates. For the 5nc10Y callable CMS spread, the reported LGM benchmark produces an NPV of 230 814 EUR, which we report alongside the QG-SV value in Table 7.10.

Table 7.10: LGM vs. QG-SV ($d = 2$) for callable CMS products. The LGM valuation evaluates CMS indices pathwise at future exercise dates where a valid trade NPV is available. The difference column is QG-SV minus LGM. EUR 10M notional.

Product	**LGM (EUR)**	**QG-SV d=2 (EUR)**	**Difference (EUR)**
Callable CMS Swap (5nc10Y)	138 373	56 203	−82 169
Callable CMS Spread (5nc10Y)	230 814	695 659	+464 845
Underlying CMS Spread Swap	*941 533*	*941 533*	*0*
Implied CMS Spread Option	*710 719*	*245 875*	*−464 844*

For the callable CMS swap (single-index, 5nc10Y), the LGM model produces 138 373 EUR versus 56 203 EUR under QG-SV ($d = 2$), a difference of −82 169 EUR in the QG-SV-minus-LGM convention. In this empirical setup the one-factor model

assigns a lower cancellation option value than the QG-SV rows for the single-index callable CMS product.

For the callable CMS spread, the LGM NPV of 230 814 EUR is substantially below the QG-SV ($d = 2$) value of 695 659 EUR, a difference of +464 845 EUR in the QG-SV-minus-LGM convention. Equivalently, the implied issuer cancellation option is much larger under LGM (710 719 EUR) than under QG-SV (245 875 EUR): the one-factor model assigns most of the underlying spread swap value to the cancellation right, leaving the holder with a small residual callable value. The theoretical mechanism is the following. In a one-factor LGM, all forward rates are driven by a single Gaussian state variable $x(t)$. Both CMS 10Y and CMS 2Y therefore respond to $x(t)$ with similar sensitivity

$$\mathrm{CMS}_{10Y}(t) \;\approx\; S_{10}(0) + \alpha_{10}\, x(t), \qquad \mathrm{CMS}_{2Y}(t) \;\approx\; S_2(0) + \alpha_2\, x(t), \tag{7.2}$$

where $\alpha_{10} \approx \alpha_2$ when the DV01 sensitivities of the two swap rates to the short rate are similar. Equation (5.13) then reduces the spread variance to a squared difference of nearly equal exposures multiplied by the common one-factor variance scale. This mechanism explains why the one-factor LGM assigns so little residual value to the callable spread relative to the QG-SV rows: with a single Gaussian driver the model under-represents the CMS-spread dynamics that the QG-SV calibration captures through additional rate factors and stochastic volatility.

Remark 7.5 (Spread-dynamics interpretation)**.** A one-factor model can generate spread volatility when the two swap-rate loadings differ, but it cannot generate independent cross-tenor rate shocks. For the tested spread-contingent payoff, the QG-SV rows therefore point to the need for at least two independent rate drivers before smile effects can be separated from missing spread dynamics. The reported LGM callable spread NPV of 230 814 EUR quantifies this gap: it is roughly one third of the QG-SV value at $d = 2$.

7.3.8 PCA-Guided Factor Study

The raw $d = 2, 3, 4$ comparison becomes more interpretable once viewed through the PCA of the calibrated $d = 4$ benchmark-tenor correlation structure [45, 41]. This PCA is computed on the calibrated model-implied benchmark-tenor correlation matrix, not on historical yield time series, and is used only as a descriptive diagnostic rather than as a pricing or factor-selection criterion. For benchmark tenors $(3.0, 1.5, 0.6, 0.2)$ and correlation decay $\alpha = 0.05$, the first principal component explains 94.5% of the

correlation structure, the first two explain 98.4%, and the first three explain 99.6%. Table 7.11 combines these PCA shares with the existing $d = 2, 3, 4$ pricing results.

Table 7.11: PCA-guided factor study. The PCA shares are computed from the calibrated $d = 4$ benchmark-tenor correlation matrix. The price columns are the existing $d = 2, 3, 4$ QG-SV results, so the table is interpretive rather than a literal no-recalibration PCA truncation. Prices in EUR.

Factor Proxy	PCA Modes	Cum. Var.	Interpretation	Swpn 5Y×10Y	CMS 10Y	Call. Spread
$d = 2$	PC1 + PC2	98.4%	Level + slope	727 866	655 931	695 659
$d = 3$	PC1 + PC2 + PC3	99.6%	Level + slope + curvature	738 266	655 931	670 090
$d = 4$	PC1 + PC2 + PC3 + PC4	100.0%	Full benchmark-tenor span	729 979	655 931	721 340

Moving from $d = 2$ to $d = 4$ is not about recovering a large block of missing covariance variance: the first two principal components already account for 98.4% of the calibrated correlation structure. Eigenvalue shares are not pricing sensitivities, however. The $d = 3$ row roughly adds a curvature mode and produces the highest 5Y×10Y swaption value but the lowest callable CMS spread value; the $d = 4$ row adds only a 0.4% residual mode in variance terms but moves the callable CMS spread back above the $d = 2$ value. Because the rows were separately re-optimised and not accompanied by independent randomised Sobol-shift stability checks, these differences are PCA-guided sensitivity diagnostics rather than resolved factor-count rankings.

Observation 7.6 (Variance explained is not a pricing criterion)**.** Table 7.11 does not justify the conclusion that "two factors are enough". Variance explained is a poor sufficient statistic for Bermudan pricing: low-variance residual modes can coincide with material callable value changes once calibration and exercise boundaries interact.

Table 7.12: Per-tenor surface calibration RMSE (bp normal vol) across $d \in \{1, \ldots, 4\}$. The calibration target is the 2026-03-17 EUR smile grid with six expiries, three tenors, and strike wings around ATM. The lowest reported overall RMSE is obtained at $d = 4$ after joint LM/Tikhonov polishing with corrected multi-curve floating-leg schedules.

d	Overall	2Y tenor	5Y tenor	10Y tenor
1	9.33	9.42	3.35	12.70
2	8.96	11.12	4.02	10.06
3	7.16	7.53	3.66	9.14
4	5.27	5.39	2.53	6.92

The surface calibration is consistent with this PCA interpretation. The one-factor specification remains too restrictive for the multi-tenor surface, while the main

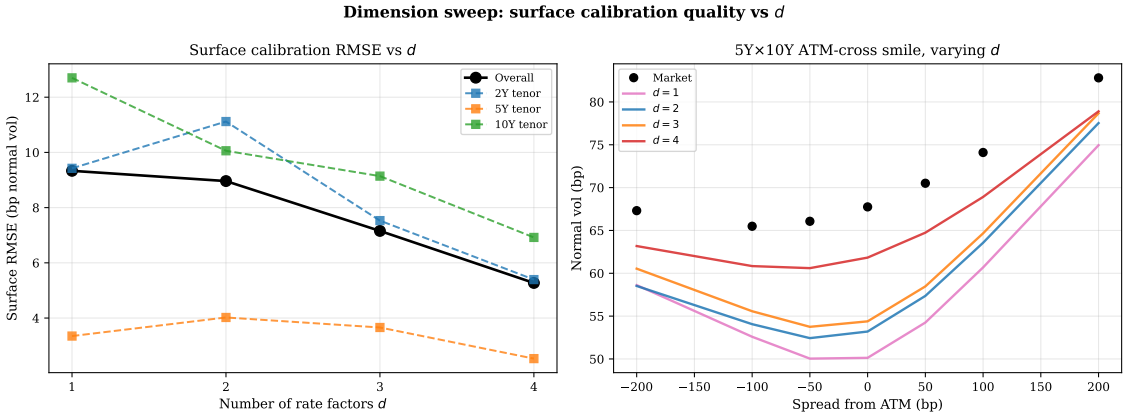

Figure 7.7: Surface calibration quality as a function of rate-factor dimension $d \in \{1, \ldots, 4\}$ using the corrected semi-analytical bootstrap followed by the joint Levenberg–Marquardt/Tikhonov polish and the corrected multi-curve floating-leg schedule. Left panel: per-tenor and overall RMSE across the 2026-03-17 EUR smile surface (six expiries $\times$ three tenors, with strike wings around ATM). The reported local-optimum fit improves from 9.33 bp for $d = 1$ to 8.96 bp for $d = 2$, reaches 7.16 bp at $d = 3$, and attains the lowest tested value of 5.27 bp at $d = 4$. Right panel: the $5Y \times 10Y$ ATM-cross smile overlaid across dimensions.

low-dimensional improvement occurs at $d = 3$, which adds a curvature mode and reduces aggregate RMSE to 7.16 bp. The lowest reported fit over the tested local candidates is obtained at $d = 4$, with a 5.27 bp overall RMSE and a 2.53 bp fit on the 5Y tenor. We therefore retain $d = 4$ as the reference surface-calibration specification, while treating the $d = 2/3/4$ product matrix as the sensitivity analysis rather than proof of an optimal factor count. The complete market-versus-model overlay across dimensions is reproduced in Appendix B.2; Figure 7.8 shows the $d = 4$ grid used as the in-text reference.

7.3.9 Full Product Comparison Matrix

Table 7.13 consolidates all pricing results into a product-by-model matrix. The callable CMS rows use the standard 5nc10Y specifications in Table 7.1: CMS 10Y versus fixed 2.7% for the single-index trade and $\max(\text{CMS}_{10Y} - \text{CMS}_{2Y} + 10\text{ bp}, 0)$ versus fixed 0% for the spread trade. The progressive decomposition results (Section 7.3.3) use a different stress specification (3× gearing, 1nc10Y callable) and are therefore reported separately.

The matrix reveals several structural patterns. First, the plain CMS control rows are identical within the analytical replication setup: the LinearTSR and Brigo–Mercurio pricing formulae use the initial curve, market swaption surface and fixed spread-correlation input directly [11, 29]. Second, European swaption prices show modest model dependence across QG-SV dimensions, with the 5Y×10Y anchor ranging from 727 K to 774 K. This dispersion is meaningful but far smaller than the model gap on spread-contingent callable products.

Third, callable CMS products exhibit the largest model sensitivity. The callable CMS spread ranges from 670 K to 721 K across the reported QG-SV dimensions, while the reported LGM row is 230 814 EUR — roughly one third of the QG-SV values. Within the QG-SV family, the non-monotone behaviour reflects the interaction between factor dimensionality and the calibration landscape.

The G2++ CMS entries are blank not because a two-factor Gaussian model is incapable of generating spread dynamics in principle, but because this thesis does not implement a like-for-like CMS coupon pricer in the G2++ framework.

7.4 Semi-Analytical vs. Monte Carlo Gap

As discussed in Remark 4.3, using semi-analytical (SA) calibration to determine model parameters and then employing those parameters in a full Monte Carlo

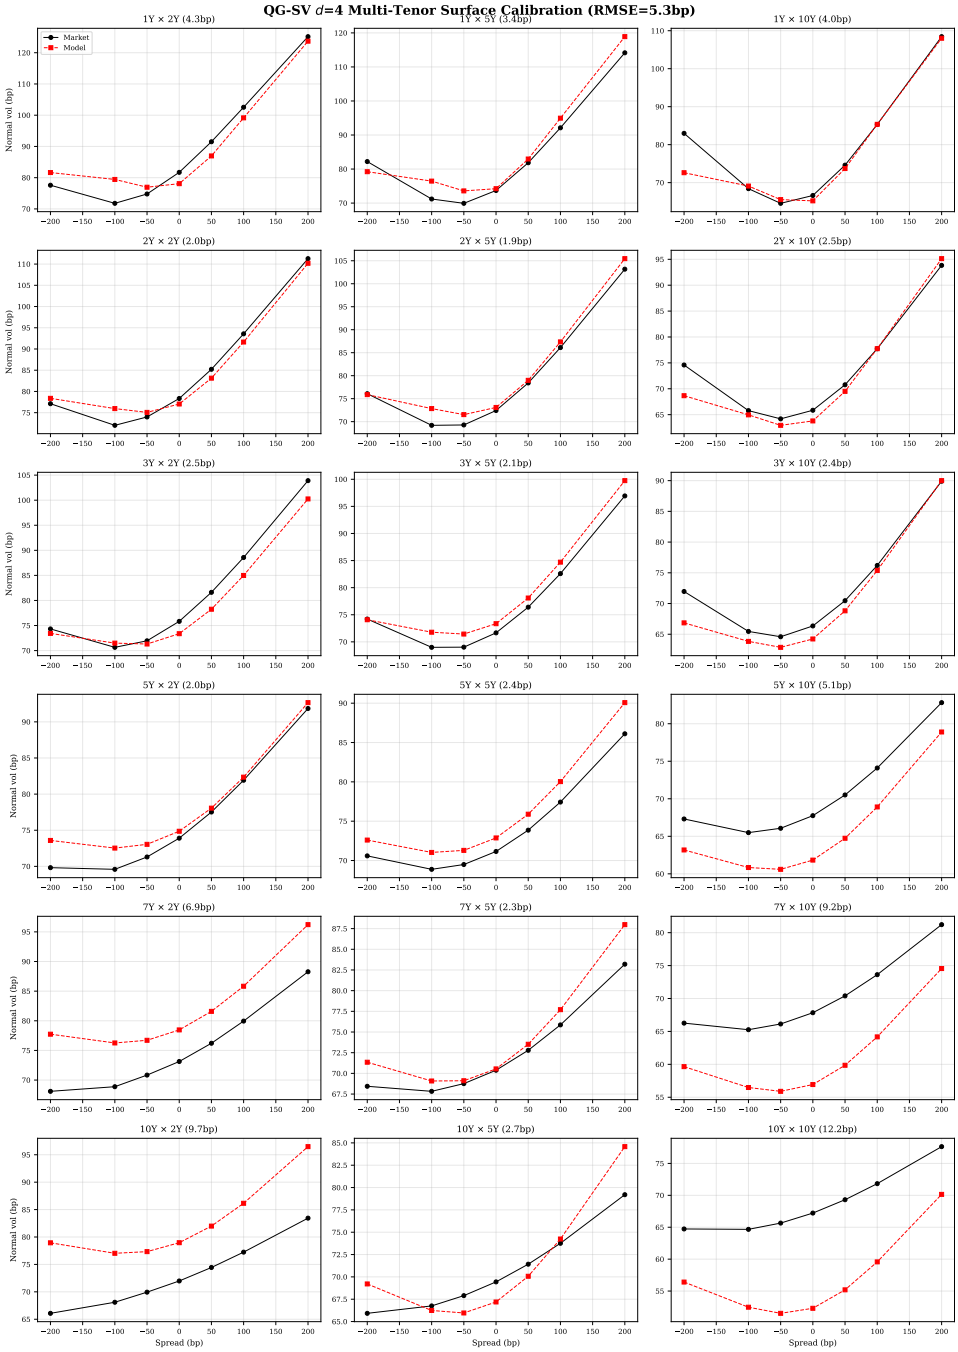

Figure 7.8: Multi-tenor surface calibration for the QG-SV $d = 4$ specification. Each panel shows the market (black circles) versus model (red squares) implied normal volatility for a given expiry–tenor combination. The calibration targets three swap tenors (2Y, 5Y, 10Y) across six expiries, with an overall RMSE of 5.27 bp under the joint LM/Tikhonov calibration workflow with explicit multi-curve floating-leg schedules.

Table 7.13: CMS product pricing matrix (EUR, 10M notional). Callable CMS products use the standard 5nc10Y specifications of Table 7.1. Plain CMS products use the same analytical pricing controls across model columns. G2++ prices selected swaptions via the strike-basket multi-curve FDM setup of Table 7.3; CMS product rows are omitted because this thesis does not report a like-for-like CMS coupon or callable-CMS pricer for the G2++ setup ("—"); d=3 and d=4 Bermudan swaption rows are not included in the reported product matrix.

Product	**LGM**	**QG-SV** d=2	**QG-SV** d=3	**QG-SV** d=4	**G2++**
Analytical CMS controls					
CMS 10Y Swap	655 931	655 931	655 931	655 931	—
CMS Spread (10Y–2Y)	941 533	941 533	941 533	941 533	—
Capped CMS 10Y	474 350	474 350	474 350	474 350	—
European swaptions (MC-priced)					
Swpn 1Y×10Y	407 868	388 556	393 160	400 955	345 883[†]
Swpn 5Y×10Y	732 340	727 866	738 266	729 979	702 341[†]
Swpn 10Y×10Y	798 458	802 665	802 737	800 406	788 650[†]
Bermudan swaption (exercise value)					
Berm. 5Y×10Y	838 655	940 619	—	—	774 584[†]
Callable CMS products (5nc10Y, model-dependent)					
Call. CMS Swap	138 373	56 203	45 462	33 850	—
Call. CMS Spread	230 814	695 659	670 090	721 340	—

[†]G2++ swaption NPVs are diagnostic strike-basket finite-difference prices under a separate constant-parameter calibration and should be read through the calibration-method caveat in Table 7.3.

(MC) simulation can produce biased prices when the frozen-sensitivity approximation changes the effective variance target. We therefore use the SA formula for diagnostics, not for setting the parameters in the MC valuation tables.

7.4.1 Mechanism

The SA swaption pricing formula in the QG-SV model relies on a *frozen-sensitivity approximation*: the state-dependent swap-rate loading $\zeta(t, X_t)$ is replaced by its time-zero proxy while the variance factor is handled by the displaced-Heston approximation. This can underestimate the effective variance experienced by MC paths when z_t and $\zeta(t, X_t)^2$ co-move positively, as formalised in Proposition A.6. Since convex option values are generally increasing in effective variance, such an underestimation can lead a calibration routine to choose larger volatility or displacement parameters than would be chosen under the full MC objective. This direction was observed in the empirical diagnostics, but it is not a general theorem about every SA calibration problem.

7.4.2 Empirical Handling in the Valuation Runs

The empirical valuation results use MC-calibrated parameters for all MC-priced instruments. Reusing SA-calibrated parameters would mix two effective calibration objectives: SA calibration freezes the swap-rate sensitivity, whereas downstream Bermudan valuation simulates the full variance process. Diagnostics on the valuation data showed this transfer to be unstable enough to obscure the market comparison, so the final tables avoid SA→MC repricing rows and use the SA approximation only for fast surface diagnostics such as Figure 7.7.

7.4.3 Implications for Practice

For callable CMS products requiring Bermudan exercise valuation, calibrating with the same MC pricing method avoids this additional approximation mismatch. The callable-pricing results therefore use MC bootstrap calibration (Algorithm 1) despite its higher computational cost.

7.5 Computational Considerations

The computational cost of stochastic-volatility models matters for repeated calibration and pricing, but absolute wall-clock timings are implementation- and hardware-

dependent. The relevant conclusion for this thesis is therefore algorithmic rather than system-specific.

The dominant cost is the Monte Carlo bootstrap calibration. Each objective evaluation requires a full simulation of the QG-SV state process, a valuation of the calibration instruments, and an implied-volatility inversion. The cost grows with the number of expiry buckets, paths, time steps and model factors, and it also increases with the number of free parameters per bucket. Higher-dimensional factor specifications are therefore more expensive not only because the state vector is larger, but also because the local optimisation problem has more directions in which to search.

The calibrated model can be reused efficiently once the piecewise parameter term structures have been fixed. This separation between calibration and repricing is important methodologically: calibration is the expensive batch task, whereas pricing with fixed parameters isolates the valuation impact of a given calibrated model. The product matrices in this chapter follow this workflow by reusing saved parameter term structures for the frozen repricing comparisons.

Semi-analytical swaption pricing remains useful as a diagnostic and warm-start tool because it avoids path simulation and gives fast residual evaluations. However, as documented in Section 7.4, its frozen-sensitivity approximation changes the calibration target. The computational advantage therefore does not by itself justify using semi-analytical parameters as inputs to full Monte Carlo callable-CMS valuation.

Remark 7.7 (Scalability Considerations)**.** The computational costs reported here reflect a research methodology without production-level optimizations. Several standard avenues for acceleration are available: parallelization across Sobol paths, adjoint algorithmic differentiation (AAD) for gradient-based calibration replacing finite-difference Jacobians, and reduced path counts with control variates or multilevel estimators; see [24, 23, 5] for Monte Carlo acceleration techniques. Quantifying those speedups for the present framework is outside the scope of the thesis.

7.6 Discussion

The results support three observations for the 2026-03-17 EUR case study. First, the analytical CMS control rows are identical by construction, while callable products inherit the model's joint forward-rate distribution. Second, smile features are economically active but not monotone: local-volatility skew and stochastic-volatility kurtosis move the issuer option in opposite directions on the stress-amplified trade,

consistent with standard smile-risk intuition [35, 30, 5]. Third, the tested spread-contingent callable products are sensitive to factor structure, because they combine CMS convexity, spread correlation and Bermudan exercise [63, 62, 31].

The most controlled empirical evidence is the nested Gaussian/QG+Skew/QG-SV comparison in Table 7.9, where model classes, calibration set, numerical method and downstream pricing method are shared. The G2++ and factor-dimension comparisons are deliberately less conclusive: G2++ uses a different calibration/pricing setup, while the $d = 2, 3, 4$ QG-SV rows are separately re-optimised and therefore mix factor count with calibration difficulty. These comparisons are useful mechanism checks, but they do not support a definitive equal-information ranking of all model families.

7.6.1 Limitations and Caveats

The limitations are correspondingly important. No liquid market prices are available for the callable CMS products, so the analysis measures model dispersion rather than pricing accuracy. The reported QG-SV valuations do not include trade-level Monte Carlo standard errors; with quasi-random Sobol paths, repeated randomised shifts or independent seeds would be needed to make the smaller $d = 2/3/4$ differences statistically resolved [67, 24]. The G2++ benchmark is constant-parameter and calibrated to a different instrument set, so it is qualitative. Finally, the calibration targets spot European swaption smiles, whereas Bermudan exercise depends on forward smiles and joint future swap-rate distributions. The chapter therefore supports mechanism identification and model-risk diagnostics, not validated market prices.

8 Conclusion

This thesis developed a multi-curve pricing framework for CMS instruments under multi-factor quasi-Gaussian stochastic volatility. The empirical study compared Gaussian, skewed quasi-Gaussian and QG-SV specifications on European swaptions, CMS spread options and callable CMS spread swaps, with emphasis on how factor structure, smile dynamics and numerical calibration consistency affect valuation.

Here d denotes the number of rate factors, and the quoted EUR NPVs use the EUR 10M notional of Table 7.1.

8.1 Summary of Findings

Under a common multi-curve setup, factor structure, skew and stochastic volatility materially affect callable CMS prices, but the effect is product-specific and not monotone in model complexity. The cross-family fixed-strike benchmark shows that the strike-basket G2++ diagnostic stays close to the vanilla swaption anchors but does not reproduce the larger QG-SV 5Y×10Y exercise premium, while the more controlled nested-QG experiment shows why this cannot be reduced to a simple ranking. Across Gaussian, QG+Skew and QG-SV variants, calibration RMSE improves from 6.87 to 3.50 to 2.00 bp and the 5Y×10Y European anchor remains close, but callable CMS swap and callable CMS spread values move differently.

The factor-dimension study further indicates that PCA variance explained is not a sufficient pricing criterion for Bermudan products. Although the first two principal components of the calibrated four-factor benchmark-tenor correlation matrix explain 98.4% of the model-implied variance, the callable CMS spread moves from 695 659 EUR at $d = 2$ to 670 090 EUR at $d = 3$ and 721 340 EUR at $d = 4$. Because these rows were separately calibrated and do not include repeated randomised Sobol shifts, they should be read as factor-sensitivity diagnostics rather than statistically resolved price rankings. Low-variance residual modes may still matter once calibration and exercise interact.

The stress-amplified 1nc10Y decomposition trade identifies one mechanism. For that trade, skew raises callable NPV by 4.2% relative to the Gaussian baseline, while stochastic volatility partly offsets it (−2.9%). The reported LGM benchmark prices the 5nc10Y callable CMS spread at 230 814 EUR, roughly one third of the QG-SV values (670–721 K), which is consistent with the one-factor spread-dynamics argument: a single Gaussian driver under-represents the cross-tenor decorrelation that the multi-factor QG-SV calibration captures. This is a case-study result, not a

statement that every multi-factor Gaussian model would behave similarly.

Semi-analytical calibration remains useful for fast surface diagnostics, but the frozen-sensitivity approximation changes the effective calibration target. In the Monte Carlo callable-pricing setting studied here, calibrating and pricing with numerically consistent methods is therefore methodologically preferable to transferring semi-analytical parameters directly into simulation.

8.2 Practical Implications

For plain CMS coupons and non-callable CMS spreads, the tested model choice is largely irrelevant because pricing is driven by analytical replication from the initial curve and swaption surface. The dynamic model matters once Bermudan exercise or spread dynamics must be generated inside the pricing model.

For single-index callable CMS products, the evidence suggests material but manageable model dependence. For callable CMS spreads, the case-study message is sharper within the QG-SV family: factor dimension and smile calibration interact with the spread option value, and at least two independent rate drivers are needed before smile effects can be meaningfully separated from missing spread dynamics.

The stress-amplified decomposition suggests that local-volatility skew captures a large correction away from Gaussian pricing, while Table 7.9 shows that stochastic volatility can still matter once the full callable portfolio is repriced under a common full-smile calibration family. The practical message is not that QG-SV is universally preferred, but that skew and stochastic volatility are economically relevant model-risk drivers for the tested callable CMS portfolio.

8.3 Limitations

The conclusions are bounded by the design of the numerical experiment.

The main limitation is scope. All calibration and pricing results are based on one EUR market date and a narrow product set. The magnitudes reported in Tables 7.3–7.11 should therefore be read as case-study evidence, not as universal constants.

Benchmark dependence is the second limitation. Table 7.3 combines different calibration methods, and Table 7.11 is PCA-guided rather than a literal no-recalibration truncation experiment. These comparisons are informative, but they support direction and mechanism more strongly than fine-grained model ranking.

Monte Carlo uncertainty and optimisation risk remain present. The reported

output does not report trade-level standard errors, and the non-monotone $d = 2/3/4$ results show that local minima can interact with economic interpretation even when the severe 7Y anomaly is absent. The dimension results should therefore be read as sensitivity analysis, not as a statistically resolved ordering of factor counts, and the conclusions are stronger at the level of economic materiality than at the level of single-price precision.

The calibration targets spot swaption smiles only. We did not perform forward-smile calibration, although Bermudan exercise depends on forward distributions. Likewise, the valuations presented here are clean prices and do not propagate the model differences into exposure or XVA calculations.

8.4 Future Work

Several extensions would address these limitations. A first extension would be to repeat Tables 7.3 and 7.9 on additional market dates and, ideally, additional currencies. This would test whether the present ranking of model effects is persistent or largely specific to the EUR surface studied here.

A second extension would be a more like-for-like Gaussian benchmark. A richer piecewise Gaussian two-factor specification together with a callable CMS pricing method would separate factor structure from smile effects more cleanly than the strike-targeted constant-parameter G2++ benchmark used here.

Although we extended the analysis to $d = 4$ factor specifications (Section 7.3.8), a systematic investigation of factor selection criteria, including information-theoretic approaches and cross-validation against held-out swaption data, would provide useful guidance for practitioners. The non-monotone dependence of callable prices on dimension suggests that the optimal factor count is product-specific and calibration-dependent, and deserves further study.

The hierarchical calibration procedure of [5], which sequentially calibrates to swaptions of increasing expiry while preserving the fit to shorter expiries, would enable forward smile calibration. This is particularly relevant for Bermudan CMS spread swaps, where the exercise boundary depends on the model-implied distribution of swap rates at each exercise date, a distribution governed by the forward rather than the spot smile.

The QG-LV model, which admits more tractable semi-analytical pricing than the full QG-SV specification, is a natural candidate for use as a control variate in Monte Carlo simulation. By computing the QG-SV minus QG-LV price difference via simulation while using the semi-analytical QG-LV price as the baseline, one could

potentially achieve large variance reduction and bring down the cost of the Monte Carlo bootstrap calibration.

Incorporating the QG-SV model into an XVA exposure framework would enable analysis of how model complexity affects credit and funding adjustments for CMS portfolios. Since XVA calculations depend on the full distribution of future portfolio values, and particularly on the tails of the swap rate distribution that are most sensitive to skew and smile, the model effects may be amplified relative to the clean price analysis presented here.

The pricing framework developed in this thesis could be extended to other CMS-linked exotic products, including target accrual redemption notes (TARNs), snowball coupons, and range accruals on CMS spreads. These products may be even more sensitive to distributional assumptions than the callable CMS spread swap, and the progressive complexity methodology could provide useful insight into model risk for these structures.

References

[1] Ferdinando Ametrano and Marco Bianchetti. Everything you always wanted to know about multiple interest rate curve bootstrapping but were afraid to ask. *SSRN Electronic Journal*, 2013. Available at SSRN 2219548.

[2] Leif B.G. Andersen and Jesper Andreasen. Volatility skews and extensions of the LIBOR market model. *Applied Mathematical Finance*, 7(1):1–32, 2000.

[3] Leif B.G. Andersen and Mark Broadie. Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science*, 50(9): 1222–1234, 2004.

[4] Leif B.G. Andersen and Vladimir V. Piterbarg. *Interest Rate Modeling. Volume 1: Foundations and Vanilla Models*. Atlantic Financial Press, 2010.

[5] Leif B.G. Andersen and Vladimir V. Piterbarg. *Interest Rate Modeling. Volume 2: Term Structure Models*. Atlantic Financial Press, 2010.

[6] Dominique Bang and Elias Daboussi. Overdue annuities and other forward LIBOR products. *Risk Magazine*, 2020. Cutting edge: Interest rate derivatives.

[7] Marco Bianchetti. Two curves, one price. *Risk Magazine*, 23(8):74–80, 2010.

[8] Tomas Björk and Bent Jesper Christensen. Interest rate dynamics and consistent forward rate curves. *Mathematical Finance*, 9(4):323–348, 1999.

[9] Tomas Björk and Lars Svensson. On the existence of finite dimensional realizations for nonlinear forward rate models. *Mathematical Finance*, 11(2):205–243, 2001.

[10] Alan Brace, Dariusz Gatarek, and Marek Musiela. The market model of interest rate dynamics. *Mathematical Finance*, 7(2):127–155, 1997.

[11] Damiano Brigo and Fabio Mercurio. *Interest Rate Models: Theory and Practice*. Springer, 2nd edition, 2006.

[12] Markus K. Brunnermeier. Deciphering the liquidity and credit crunch 2007–2008. *Journal of Economic Perspectives*, 23(1):77–100, 2009.

[13] Peter Carr and Dilip Madan. Option valuation using the fast Fourier transform. *Journal of Computational Finance*, 2(4):61–73, 1999.

[14] Peter Carr and Dilip Madan. Towards a theory of volatility trading. In *Volatility: New Estimation Techniques for Pricing Derivatives*, pages 417–427. Risk Books, 2001.

[15] Andrew Carverhill. When is the short rate Markovian? *Mathematical Finance*, 4(4):305–312, 1994.

[16] Simon Cedervall and Vladimir V. Piterbarg. Full implications for CMS convexity. *Asia Risk Magazine*, 2012. Cutting edge: Interest rate derivatives.

[17] Oren Cheyette. Term structure dynamics and mortgage valuation. *The Journal of Fixed Income*, 1(4):28–41, 1992. doi: 10.3905/jfi.1992.408036.

[18] John C. Cox, Jonathan E. Ingersoll, and Stephen A. Ross. A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407, 1985.

[19] William Feller. Two singular diffusion problems. *Annals of Mathematics*, 54(1): 173–182, 1951.

[20] Claudio Fontana, Giacomo Lanaro, and Agatha Murgoci. The geometry of multi-curve interest rate models. *arXiv preprint arXiv:2401.11619v2*, 2024. To appear in Mathematical Finance.

[21] Masaaki Fujii, Yasufumi Shimada, and Akihiko Takahashi. A note on construction of multiple swap curves with and without collateral. Technical Report F-154, Center for Advanced Research in Finance, University of Tokyo, 2010.

[22] Hélyette Geman, Nicole El Karoui, and Jean-Charles Rochet. Changes of numeraire, changes of probability measure and option pricing. *Journal of Applied Probability*, 32(2):443–458, 1995.

[23] Michael B. Giles. Multilevel Monte Carlo path simulation. *Operations Research*, 56(3):607–617, 2008.

[24] Paul Glasserman. *Monte Carlo Methods in Financial Engineering*. Applications of Mathematics. Springer, 2003.

[25] Paul Glasserman and Bin Yu. Number of paths versus number of basis functions in American option pricing. *Annals of Applied Probability*, 14(4):2090–2119, 2004.

[26] Gene H. Golub and Charles F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, 4th edition, 2013.

[27] Zorana Grbac and Wolfgang J. Runggaldier. *Interest Rate Modeling: Post-Crisis Challenges and Approaches*. SpringerBriefs in Quantitative Finance. Springer International Publishing, 2015.

[28] István Gyöngy. Mimicking the one-dimensional marginal distributions of processes having an Itô differential. *Probability Theory and Related Fields*, 71 (4):501–516, 1986.

[29] Patrick S. Hagan. Convexity conundrums: pricing CMS swaps, caps, and floors. *Wilmott Magazine*, 5:38–44, 2003.

[30] Patrick S. Hagan, Deep Kumar, Andrew S. Lesniewski, and Diana E. Woodward. Managing smile risk. *Wilmott Magazine*, pages 84–108, 2002.

[31] Patrick S. Hagan, Andrew S. Lesniewski, Georgios E. Skoufis, and Diana E. Woodward. CMS spread options. *Quantitative Finance*, 21(9):1539–1554, 2021. doi: 10.1080/14697688.2021.1917122.

[32] David Heath, Robert Jarrow, and Andrew Morton. Bond pricing and the term structure of interest rates: a new methodology for contingent claims valuation. *Econometrica*, 60(1):77–105, 1992.

[33] Marc Henrard. *Interest Rate Modelling in the Multi-Curve Framework: Foundations, Evolution and Implementation*. Applied Quantitative Finance. Palgrave Macmillan, 2014.

[34] Marc Henrard. A quant perspective on IBOR fallback proposals. *SSRN Electronic Journal*, 2018. Available at SSRN: https://ssrn.com/abstract=3226183.

[35] Steven L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6 (2):327–343, 1993.

[36] Thomas S. Y. Ho and Sang Bin Lee. Term structure movements and pricing interest rate contingent claims. *Journal of Finance*, 41(5):1011–1029, 1986.

[37] John Hull and Alan White. Pricing interest-rate-derivative securities. *Review of Financial Studies*, 3(4):573–592, 1990.

[38] Chris Hunter and Peter Jäckel. Replication of CMS derivatives. *Wilmott Magazine*, 2003.

[39] Farshid Jamshidian. An exact bond option formula. *Journal of Finance*, 44(1): 205–209, 1989.

[40] Farshid Jamshidian. LIBOR and swap market models and measures. *Finance and Stochastics*, 1(4):293–330, 1997.

[41] Ian T. Jolliffe. *Principal Component Analysis*. Springer, 2nd edition, 2002.

[42] Mark S. Joshi and Riccardo Rebonato. A stochastic-volatility, displaced-diffusion extension of the LIBOR market model. *Quantitative Finance*, 3: 458–469, 2003.

[43] Nikolaos Karouzakis, John Hatgioannides, and Kostas Andriosopoulos. Convexity adjustment for constant maturity swaps in a multi-curve framework. *Annals of Operations Research*, 266:433–456, 2017. doi: 10.1007/s10479-017-2430-6.

[44] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944.

[45] Robert Litterman and José Scheinkman. Common factors affecting bond returns. *Journal of Fixed Income*, 1(1):54–61, 1991.

[46] Francis A. Longstaff and Eduardo S. Schwartz. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14 (1):113–147, 2001.

[47] Roger Lord, Remmert Koekkoek, and Dick Van Dijk. A comparison of biased simulation schemes for stochastic volatility models. *Quantitative Finance*, 10 (2):177–194, 2010.

[48] Andrey Lyashenko and Fabio Mercurio. Looking forward to backward-looking rates. *Risk Magazine*, 2019.

[49] Donald W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11 (2):431–441, 1963.

[50] Fabio Mercurio. Interest rates and the credit crunch: New formulas and market models. *Bloomberg Portfolio Research Paper*, 2009. Available at SSRN: https://ssrn.com/abstract=1331600.

[51] Fabio Mercurio. A simple multi-curve model for pricing sofr futures and other derivatives. *SSRN working paper*, 2018.

[52] Kristian R. Miltersen, Klaus Sandmann, and Dieter Sondermann. Closed form solutions for term structure derivatives with log-normal interest rates. *Journal of Finance*, 52(1):409–430, 1997.

[53] John A. Nelder and Roger Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, 1965.

[54] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, 2nd edition, 2006.

[55] Open Source Risk Engine contributors. Open Source Risk Engine. URL https://www.opensourcerisk.org/. Source code: https://github.com/OpenSourceRisk/Engine; Modified BSD License.

[56] Antoon Pelsser. Mathematical foundation of convexity correction. *Quantitative Finance*, 3(1):59–65, 2003.

[57] Goran Peskir and Albert N. Shiryaev. *Optimal Stopping and Free-Boundary Problems*. Lectures in Mathematics ETH Zürich. Birkhäuser, 2006.

[58] Vladimir V. Piterbarg. Stochastic volatility model with time-dependent skew. *Applied Mathematical Finance*, 12(2):147–185, 2005.

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

[60] Vladimir V. Piterbarg. The IBOR fallback and its implications. *Risk Magazine*, 2020.

[61] Dmitry Pugachevsky. Forward CMS rate adjustment. *Risk Magazine*, March 2003.

[62] Riccardo Rebonato. *Volatility and Correlation: The Perfect Hedger and the Fox*. Wiley, 2nd edition, 2004.

[63] Riccardo Rebonato and Kenneth McKay. Correlation and hedging of callable LIBOR exotics. *Risk Magazine*, 15(10):119–123, 2002.

[64] Peter Ritchken and L. Sankarasubramanian. Volatility structures of forward rates and the dynamics of the term structure. *Mathematical Finance*, 5(1): 55–72, 1995.

[65] L. C. G. Rogers and Martin B. Haugh. Monte Carlo valuation of American options. *Mathematical Finance*, 12(3):271–286, 2002.

[66] Steven E. Shreve. *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer Finance. Springer, 2004.

[67] Ilya M. Sobol. On the distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics*, 7(4):86–112, 1967.

[68] The QuantLib contributors. QuantLib: a free/open-source library for quantitative finance. URL https://www.quantlib.org/. DOI: https://doi.org/10.5281/zenodo.1440997; source code: https://github.com/lballabio/QuantLib; license: BSD-3-Clause.

[69] Andrey N. Tikhonov and Vasiliy Y. Arsenin. *Solutions of Ill-Posed Problems*. Winston, 1977.

[70] Oldrich Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188, 1977.

[71] Servaas Willems. Backward-looking RFR caplets with the SABR model. *Risk Magazine*, 2020.

A Technical Proofs and Derivations

This appendix collects the derivations that are too long for the main text but are central to the mathematical content of the thesis. The notation is that of Chapters 2–5. Throughout, the standing assumptions stated before Chapter 2 are in force.

A.1 Collateralised Multi-Curve Pricing

Theorem A.1 (Collateralised pricing identity)**.** *Let X_T be an integrable payoff at time T. Under OIS collateralisation, its time-t value is*

$$V_t = P^d(t,T)\,\mathbb{E}^T[X_T \mid \mathcal{F}_t]\,, \qquad 0 \le t \le T, \tag{A1}$$

where $\mathbb{Q}^T$ is the T-forward measure associated with the numeraire $P^d(\cdot,T)$. The density process is

$$L_u^T = \left.\frac{d\mathbb{Q}^T}{d\mathbb{Q}}\right|_{\mathcal{F}_u} = \frac{B_0^d}{P^d(0,T)}\,\frac{P^d(u,T)}{B_u^d}, \qquad 0 \le u \le T,$$

which is a true martingale by Assumption 0.1. Equivalently,

$$\frac{V_t}{B_t^d} = \mathbb{E}^{\mathbb{Q}}\left[\frac{X_T}{B_T^d} \mid \mathcal{F}_t\right], \tag{A2}$$

where B^d is the OIS money-market account. This is the standard martingale-pricing and change-of-numeraire identity [22, 66].

Proof. By Assumptions 0.1 and 0.4, discounted collateralised asset prices are true $\mathbb{Q}$-martingales and X_T/B_T^d is integrable. Hence

$$V_t = B_t^d \mathbb{E}^{\mathbb{Q}}\left[X_T/B_T^d \mid \mathcal{F}_t\right].$$

At maturity the density above reduces to the usual change-of-numeraire density [22]

$$\left.\frac{d\mathbb{Q}^T}{d\mathbb{Q}}\right|_{\mathcal{F}_T} = \frac{B_T^{d,-1}P^d(T,T)}{B_0^{d,-1}P^d(0,T)} = \frac{B_0^d}{B_T^d P^d(0,T)}.$$

Applying Bayes' formula conditionally on $\mathcal{F}_t$ gives

$$\mathbb{E}^{\mathbb{Q}}\left[X_T/B_T^d \mid \mathcal{F}_t\right] = \frac{P^d(t,T)}{B_t^d}\mathbb{E}^T[X_T \mid \mathcal{F}_t],$$

which proves (A1). □

A.2 Quasi-Gaussian Bond Pricing

The exponential-affine representation below is the standard QG bond-pricing formula of Andersen–Piterbarg's finite-dimensional term-structure construction [5].

Proposition A.2 (Exponential-affine bond representation)**.** *In the d-factor QG-SV model of Definition 4.5, under Assumption 0.4, define*

$$G_k(t,T) = \int_t^T \exp\Big(-\int_t^u \kappa_k(s)\,ds\Big)\,du.$$

Then the OIS zero-coupon bond price is

$$P(t,T) = \frac{P(0,T)}{P(0,t)} \exp\left(-\sum_{k=1}^{d} G_k(t,T)x_k(t) - \frac{1}{2}\sum_{k,l=1}^{d} G_k(t,T)G_l(t,T)y_{kl}(t)\right). \quad \text{(A3)}$$

Proof. Let the right-hand side of (A3) be denoted by $\widehat{P}(t,T)$. It is enough to show that the OIS-discounted process $\widehat{P}(t,T)/B_t^d$ has zero drift, that the resulting stochastic integral is a true martingale under Assumption 0.4, and that $\widehat{P}(T,T) = 1$. The terminal condition follows from $G_k(T,T) = 0$.

Write

$$\log \widehat{P}(t,T) = \log P(0,T) - \log P(0,t) - G(t,T)^\top x(t) - \frac{1}{2}G(t,T)^\top y(t)G(t,T).$$

Applying Itô's formula and using

$$dG_k(t,T) = \Big(\kappa_k(t)G_k(t,T) - 1\Big)dt$$

together with the QG-SV state equations gives the diffusion

$$\frac{d\widehat{P}(t,T)}{\widehat{P}(t,T)} = r(t)\,dt - \sum_{k=1}^{d} G_k(t,T)[\sigma_x(t,x,y)]_{k,\cdot}\sqrt{z(t)}\,dW^x(t).$$

The key cancellation is the covariance part of this calculation. Since $d\langle x\rangle_t = z(t)\sigma_x\sigma_x^\top dt$, the dy-term in $d\log\widehat{P}$ contributes

$$-\frac{1}{2}z(t)\,G(t,T)^\top\sigma_x\sigma_x^\top G(t,T)\,dt,$$

while exponentiating $\log \widehat{P}$ contributes the Itô correction

$$\frac{1}{2} d\langle G^\top x \rangle_t = \frac{1}{2} z(t)\, G(t,T)^\top \sigma_x \sigma_x^\top G(t,T)\, dt.$$

These two terms cancel exactly; the remaining finite-variation terms combine to the short-rate drift $r(t)dt$. Therefore

$$d\left(\frac{\widehat{P}(t,T)}{B_t^d}\right) = -\frac{\widehat{P}(t,T)}{B_t^d} \sum_{k=1}^{d} G_k(t,T)[\sigma_x(t,x,y)]_{k,\cdot} \sqrt{z(t)}\, dW^x(t),$$

which has zero drift. Assumption 0.4 requires the stochastic exponential generated by this diffusion to be a true martingale. Hence $\widehat{P}(t,T)/B_t^d = \mathbb{E}^{\mathbb{Q}}[1/B_T^d \mid \mathcal{F}_t]$. By the martingale pricing theorem [66], equivalently by uniqueness of the arbitrage-free discounted price process with terminal payoff 1 at T, $\widehat{P}(t,T)$ coincides with the OIS bond price $P(t,T)$. □

A.3 Swap-Rate Dynamics Under the Annuity Measure

The following derivation combines the QG bond representation with the change-of-numeraire theorem for the swap annuity numeraire [22, 5].

Proposition A.3 (Gradient of the par swap rate)**.** *In the equivalent single-curve notation $S^{a,b}(t) = (P(t,T_a) - P(t,T_b))/A^{a,b}(t)$, the derivative of the swap rate with respect to the QG factor x_k is*

$$\begin{aligned}\nabla_{x_k} S^{a,b}(t) = -\frac{1}{A^{a,b}(t)} \Bigg(& G_k(t,T_a)P(t,T_a) - G_k(t,T_b)P(t,T_b) \\ & - S^{a,b}(t) \sum_{i=a+1}^{b} \tau_i G_k(t,T_i) P(t,T_i) \Bigg).\end{aligned} \tag{A4}$$

For a multi-curve par rate with deterministic basis cashflows, the same quotient rule applies after replacing $P(t,T_a) - P(t,T_b)$ by the projected floating-leg present value; additional deterministic-basis cashflow terms then enter through their discount-bond sensitivities.

Proof. Let $N(t) = P(t,T_a) - P(t,T_b)$ and $A(t) = \sum_{i=a+1}^{b} \tau_i P(t,T_i)$. From (A3),

$$\partial_{x_k} P(t,T) = -G_k(t,T) P(t,T).$$

The quotient rule gives

$$\partial_{x_k} S = \frac{(\partial_{x_k} N)A - N(\partial_{x_k} A)}{A^2}.$$

Substituting

$$\partial_{x_k} N = -G_k(t, T_a)P(t, T_a) + G_k(t, T_b)P(t, T_b), \quad \partial_{x_k} A = -\sum_i \tau_i G_k(t, T_i)P(t, T_i),$$

and using $N = SA$ yields (A4). □

Proposition A.4 (Annuity-measure swap-rate SDE)**.** *Under Assumptions 0.2 and 0.4, so that deterministic basis spreads do not add stochastic state variables to the par-rate numerator and the annuity-measure value process is a true martingale, let $S^{a,b}(t) = \Pi_f^{a,b}(t)/A^{a,b}(t)$, where $\Pi_f^{a,b}$ is the projected floating-leg present value of Definition 3.1. The swap rate satisfies under the annuity measure $\mathbb{Q}^A$*

$$dS^{a,b}(t) = \sum_{k=1}^{d} \nabla_{x_k} S^{a,b}(t)[\sigma_x(t, x, y)]_{k,\cdot} \sqrt{z(t)}\, dW^A(t), \tag{A5}$$

where W^A is a d-dimensional Brownian motion with independent components, consistent with Definition 4.5, and

$$\nabla_{x_k} S^{a,b}(t) = \frac{(\partial_{x_k} \Pi_f^{a,b})(t) A^{a,b}(t) - \Pi_f^{a,b}(t)(\partial_{x_k} A^{a,b})(t)}{\left(A^{a,b}(t)\right)^2}.$$

When $L^f = L^d$, this quotient derivative reduces to (A4). With deterministic basis cashflows it additionally contains the derivatives of the discounted basis terms, but no new Brownian drivers.

Proof. Under Assumption 0.2, the basis adjustment is deterministic and the stochastic part of the par-rate numerator is a function of OIS bond prices. The ratio $S^{a,b}(t) = \Pi_f^{a,b}(t)/A^{a,b}(t)$ is therefore the value, in annuity units, of the corresponding floating-minus-fixed bond portfolio, and is a martingale under the annuity measure by change of numeraire and Assumption 0.4. Applying Itô's formula to $S^{a,b}(t, x(t), y(t))$, the annuity-measure drift is zero and the diffusion part is the gradient with respect to the stochastic factors multiplied by the factor diffusion. The y-variables have finite variation, so they contribute no diffusion term directly. □

A.4 Static Replication Formula

The replication identity is the Carr–Madan second-derivative representation, applied under the annuity measure with swaptions as the option instruments [14, 38].

Proposition A.5 (Carr–Madan representation for swap-rate payoffs)**.** *Let S_T be a swap rate that is a martingale under $\mathbb{Q}^A$, with forward level $F = S_0$. If $h \in C^2$ and*

$$\int_F^\infty |h''(K)|\mathbb{E}^A[(S_T - K)^+]\,dK + \int_{-\infty}^F |h''(K)|\mathbb{E}^A[(K - S_T)^+]\,dK < \infty,$$

then

$$\mathbb{E}^A[h(S_T)] = h(F) + \int_F^\infty h''(K)\mathbb{E}^A[(S_T - K)^+]\,dK + \int_{-\infty}^F h''(K)\mathbb{E}^A[(K - S_T)^+]\,dK.$$

Proof. For fixed s, split Taylor's formula with integral remainder at F

$$h(s) = h(F) + h'(F)(s - F) + \int_F^\infty h''(K)(s - K)^+ dK + \int_{-\infty}^F h''(K)(K - s)^+ dK.$$

Taking expectations is justified by the stated integrability condition and Fubini's theorem. Using $\mathbb{E}^A[S_T - F] = 0$ proves the claim. □

A.5 Analytical CMS Pricing Methods Used as Controls

This section records the analytical methods used for the non-callable CMS control rows in Chapter 7. They are included here to make clear why those rows are identical across model columns in the numerical study: the valuation inputs are the initial curve, the market swaption smile and, for spreads, an exogenous spread-correlation approximation, rather than paths generated by the QG-SV, LGM or G2++ dynamics. The equality is therefore conditional on this fixed analytical control stack and should not be read as a general model-invariance result for CMS spread payoffs.

Linear terminal swap-rate CMS pricing. For a CMS coupon depending on one swap rate $S_T = S^{a,b}(T)$, Proposition A.5 applies after changing from the payment measure to the annuity measure. The exact payment-measure expectation can be written as

$$\mathbb{E}^{T_p}[S_T] = \mathbb{E}^A[S_T \Lambda_T], \qquad \Lambda_T = \frac{P^d(T, T_p)A^{a,b}(0)}{A^{a,b}(T)P^d(0, T_p)}.$$

The LinearTSR approximation replaces the conditional annuity mapping by an affine function of the terminal swap rate,

$$\alpha_T(s) := \mathbb{E}^A[\Lambda_T \mid S_T = s] \approx a_T + b_T s.$$

The CMS expectation is then approximated by

$$\mathbb{E}^{T_p}[S_T] \approx \mathbb{E}^A[S_T(a_T + b_T S_T)].$$

More generally, capped and floored CMS coupons use $\mathbb{E}^A[\phi(S_T)(a_T + b_T S_T)]$, where ϕ is the coupon payoff function. The terms involving nonlinear functions of S_T are evaluated from the market swaption smile through the static-replication formula. Thus the dynamic interest-rate model is not used to simulate the coupon; it enters only if the product contains exercise decisions or path dependence. This is the linear-terminal-swap-rate implementation of the CMS convexity adjustment discussed by Hagan and in the replication literature [29, 38, 56, 61, 16].

Brigo–Mercurio CMS spread approximation. For a spread coupon depending on

$$Y_T = \alpha_1 S_T^{(1)} - \alpha_2 S_T^{(2)} + s,$$

one-dimensional replication is not sufficient because the value depends on the joint law of the two swap rates. The analytical control used in the empirical tables follows the Brigo–Mercurio spread approximation [11, 50]: the marginal distributions are anchored to the relevant swaption volatilities and the joint law is closed by a spread-correlation input. In the normal-volatility notation used for the EUR smile diagnostics, the corresponding first-order spread variance is

$$\sigma_Y^2(T) = \alpha_1^2\sigma_1^2(T) + \alpha_2^2\sigma_2^2(T) - 2\alpha_1\alpha_2\rho_{12}(T)\sigma_1(T)\sigma_2(T),$$

with forward level

$$F_Y = \alpha_1 F_1 - \alpha_2 F_2 + s.$$

A caplet-style spread payoff is then priced by the Bachelier expression

$$\mathbb{E}[(Y_T - K)^+] \approx (F_Y - K)N(d) + \sigma_Y(T)\sqrt{T}\, n(d), \qquad d = \frac{F_Y - K}{\sigma_Y(T)\sqrt{T}},$$

with the usual limiting interpretation when $\sigma_Y(T) = 0$. Linear spread coupons use the corresponding approximate forward expectation. The approximation therefore imports smile and correlation information directly through analytical market inputs. It is a useful non-callable control, but it should not be read as evidence that the dynamic model is irrelevant for callable CMS spreads, where the exercise strategy depends on simulated future joint distributions.

A.6 Semi-Analytical Versus Monte Carlo Calibration

The semi-analytical QG-SV approximation uses the Andersen–Piterbarg frozen swap-rate sensitivity and displaced-Heston Fourier machinery [35, 13, 5].

Proposition A.6 (Bias induced by frozen sensitivities)**.** *Let the exact annuity-measure swap-rate dynamics be*

$$dS_t = \sqrt{z_t}\,\zeta(t, X_t)\,dW_t^A,$$

where $X_t = (x(t), y(t))$, and assume the integrated variances below are integrable. Let the semi-analytical approximation replace $\zeta(t, X_t)$ by a deterministic function $\zeta_0(t)$. Then the expected difference between exact and frozen integrated variance is

$$\mathbb{E}^A\left[\int_0^T z_t\zeta(t, X_t)^2 dt\right] - \mathbb{E}^A\left[\int_0^T z_t\zeta_0(t)^2 dt\right] = \int_0^T \mathbb{E}^A\Big[z_t\{\zeta(t, X_t)^2 - \zeta_0(t)^2\}\Big]\,dt. \quad \text{(A6)}$$

If, for almost every $t \in [0, T]$, $\mathbb{E}^A[z_t] > 0$, $\zeta_0(t)^2 \le \mathbb{E}^A[\zeta(t, X_t)^2]$, and the covariance between z_t and $\zeta(t, X_t)^2$ is non-negative, the frozen approximation underestimates the exact integrated variance.

Proof. The identity follows by linearity of integration and expectation. The sign condition follows after writing

$$\mathbb{E}[z_t\{\zeta(t, X_t)^2 - \zeta_0(t)^2\}] = \mathbb{E}[z_t]\Big(\mathbb{E}[\zeta(t, X_t)^2] - \zeta_0(t)^2\Big) + \mathrm{Cov}(z_t, \zeta(t, X_t)^2)$$

and using $\mathbb{E}[z_t] > 0$. The assumptions are precisely the conditions under which freezing the state-dependent swap-rate sensitivity removes a positive convexity contribution from the variance. □

The proposition is deliberately stated as a sufficient condition on integrated variance, not as a universal option-price ordering theorem. It justifies the methodological caution used in the empirical chapters: parameters calibrated with the frozen

semi-analytical approximation are not automatically valid inputs for a full Monte Carlo pricing method unless the approximation error has been checked for the product and market under study.

A.7 Sequential Bootstrap Dependence

Proposition A.7 (Locality of calibrated buckets)**.** *Let the continuous-time QG-SV parameters be piecewise constant on* $[T_{n-1}, T_n)$*, with fixed structural functions* κ_k *and fixed tenor map* Δ_k*. A European swaption expiring at* T_n *depends on calibrated parameter buckets* $1, \dots, n$ *only. A finite-path Monte Carlo estimate of that value inherits the same parameter locality, but also carries simulation and time-discretisation error. If structural functions entering post-expiry bond loadings, such as* κ_k*, were also calibrated in later buckets, this locality would not be automatic.*

Proof. The state at T_n is the strong solution of the QG-SV SDE on $[0, T_n]$. On each subinterval $[T_{j-1}, T_j)$, the coefficients depend only on bucket j. Therefore the solution stopped at T_n is a measurable functional of the Brownian path on $[0, T_n]$ and the parameter vector $(\theta^{(1)}, \dots, \theta^{(n)})$. A European swaption payoff is obtained by applying to this state the value of the underlying swap at T_n. The bond-pricing formula shows that this payoff map contains deterministic loadings $G_k(T_n, T_i)$ from T_n to later payment dates. In the parametrisation used in this thesis these loadings are determined by the fixed structural functions κ_k, the initial curve and the payment schedule, not by later calibrated (λ, b, η) buckets. Hence later calibrated buckets do not enter the swaption value. If post-expiry structural functions such as κ_k were also calibrated, the same bond-loading formula shows why the locality statement would need to be restated. □

A.8 Finite-Dimensional Realisation

This section makes precise the finite-dimensional-realisation statement used in Chapter 4. The terminology follows the Cheyette/HJM finite-dimensional lineage [17, 64, 15, 8, 9].

Definition A.8 (Finite-dimensional realisation of an HJM curve)**.** Consider an arbitrage-free HJM discount-curve model under the collateralised risk-neutral measure. The model admits an n-dimensional finite-dimensional realisation on a horizon $[0, T^*]$ if there exist an n-dimensional Markov process X and deterministic functions

$\mathcal{P}(t,T,x)$, defined for $0 \le t \le T \le T^*$, such that

$$P^d(t,T) = \mathcal{P}(t,T,X_t)$$

for every maturity T. Equivalently, the instantaneous forward curve

$$f(t,T) = -\partial_T \log P^d(t,T)$$

evolves on the finite-dimensional manifold $\{-\partial_T \log \mathcal{P}(t,\cdot,x) : x \in \mathbb{R}^n\}$. The deterministic initial curve and deterministic parameter functions are not counted as stochastic state variables.

Proposition A.9 (FDR of the benchmark-rate QG-SV model)**.** *For fixed d, the benchmark-rate QG-SV model admits a finite-dimensional realisation of the OIS HJM discount curve with state dimension*

$$n = d + \frac{d(d+1)}{2} + 1.$$

Proof. Let

$$X_t = \Big(x_1(t), \dots, x_d(t), \{y_{kl}(t)\}_{1\le k\le l\le d}, z(t)\Big).$$

The symmetry $y_{kl} = y_{lk}$ gives $d(d+1)/2$ distinct covariance variables, so X has dimension $d + d(d+1)/2 + 1$. The dynamics in Definition 4.5 are closed in these variables: the drift and diffusion of x, the finite-variation dynamics of y, and the CIR dynamics of z depend on time, deterministic parameter functions and the current value of X_t, but not on the full forward curve as an additional stochastic input. Thus X is a finite-dimensional Markov state under the standing regularity assumptions.

Define, for $x \in \mathbb{R}^d$, symmetric y, and $z \ge 0$,

$$\mathcal{P}(t,T,x,y,z) = \frac{P^d(0,T)}{P^d(0,t)} \exp\left(-\sum_{k=1}^{d} G_k(t,T)x_k - \frac{1}{2}\sum_{k,l=1}^{d} G_k(t,T)G_l(t,T)y_{kl} \right).$$

Although z does not enter this function explicitly, it enters the state dynamics and therefore the future distribution of (x,y). Proposition A.2 shows that the arbitrage-free OIS bond price satisfies

$$P^d(t,T) = \mathcal{P}(t,T,X_t)$$

for every T. Hence the entire discount curve is generated by the finite state X_t, proving the realisation property in Definition A.8.

The associated forward curve is also finite-dimensional. Differentiating the negative log bond map,

$$-\log \mathcal{P}(t, T, X_t),$$

with respect to T gives

$$f(t,T) = f(0,T) + \sum_{k=1}^{d} \partial_T G_k(t,T) x_k(t) + \frac{1}{2} \sum_{k,l=1}^{d} \partial_T (G_k(t,T) G_l(t,T))\, y_{kl}(t),$$

where

$$\partial_T G_k(t,T) = \exp\left(- \int_t^T \kappa_k(s)\, ds \right).$$

Thus the forward curve remains on the manifold parameterised by X_t.

Finally, the construction is an HJM realisation rather than merely a curve parametrisation. From Proposition A.2,

$$\frac{dP^d(t,T)}{P^d(t,T)} = r(t)\, dt - \sum_{k=1}^{d} G_k(t,T) [\sigma_x(t, X_t)]_{k,\cdot} \sqrt{z(t)}\, dW_t^x.$$

In an HJM model the bond volatility is the maturity integral of the instantaneous forward-rate volatility. Therefore the corresponding forward-rate volatility is

$$\boldsymbol{\sigma}_f(t,T) = \sum_{k=1}^{d} \partial_T G_k(t,T) [\sigma_x(t, X_t)]_{k,\cdot} \sqrt{z(t)}.$$

The zero drift of the discounted bond established in Proposition A.2 is exactly the HJM no-arbitrage drift restriction for this volatility. Consequently the benchmark-rate QG-SV model is an arbitrage-free finite-dimensional realisation of the OIS HJM curve. □

B Calibration Output and Supplementary Diagnostics

This appendix records the numerical outputs behind the $d = 2$ Monte Carlo bootstrap calibrations used in Chapter 7. Unless stated otherwise, the numbers are taken from the Gaussian, skew-only, and full QG-SV calibrations to the 2026-03-17 EUR market data. The calibration strip consists of 1Y, 2Y, 3Y, 5Y, 7Y, and 10Y expiries into the 10Y swap tenor, each with strike spreads of −200, −100, −50, 0, +50, +100, and +200 basis points around the ATM forward.

B.1 Per-Expiry Calibration Quality

Table B1 gives the per-expiry RMS strip for the three nested variants used in Chapter 7. The ATM forward column is taken from the full QG-SV calibration and is common to the nested comparison.

Table B1: Per-expiry calibration errors for the d=2 nested Gaussian/QG+Skew/QG-SV calibrations. All errors are in basis points of normal implied volatility.

Expiry	ATM fwd (%)	MF Gaussian	QG+Skew	Full QG-SV
1Y	3.036	10.6	6.1	0.5
2Y	3.123	8.7	3.4	1.6
3Y	3.206	7.3	3.5	1.8
5Y	3.367	4.2	2.1	3.1
7Y	3.445	3.7	2.2	2.4
10Y	3.469	2.9	1.6	1.7

The reported calibration output supports the main points used in Chapter 7. Among the local-optimisation candidates, the full QG-SV model achieves the lowest overall RMSE (2.00 bp across the strip), and every expiry-RMS bucket is within 3.1 bp. Aggregate RMSE values in the main text are reported to two decimals because one-decimal values would hide the distinction between rounding and truncation; the per-expiry table is rounded to one decimal for readability. Pointwise strike residuals are larger in some wings, as shown in Section B.4. No single expiry dominates the 2026-03-17 calibration residuals; the largest residual bucket is 5Y. The aggregate bar chart is shown once in Figure 7.5; it is not repeated here.

B.2 Full Surface Calibration Grid by Dimension

The dimension sweep summarised by Figure 7.7 and Table 7.12 in Section 7.3.8 is derived from four reported multi-tenor surface calibrations ($d = 1, \ldots, 4$) on the 2026-03-17 EUR smile grid (six expiries $\times$ three tenors with strike wings around ATM). Exploratory $d = 5, \ldots, 8$ figures were also generated during development, but they are excluded from the thesis claims because the higher-dimensional local optimisation became unstable and no robust global search was performed. The reported calibrations use the corrected semi-analytical bootstrap followed by the joint LM/Tikhonov polish and explicit multi-curve floating-leg schedules. Figure B1 reproduces the complete market-versus-model overlay for the reported dimensions on the same axes, so per-cell changes as a function of d can be inspected directly. Each panel title shows the per-cell RMSE for every reported dimension. The main reported pattern is aggregate improvement over the tested range, with $d = 4$ providing the lowest local-optimum overall and per-tenor surface RMSE; this remains subject to the local-optimisation caveat stated in Chapter 7.

B.3 Piecewise Parameters of the Full QG-SV Run

Table B2 reports the piecewise parameters recovered from the full QG-SV calibration. These are the same flattened interval values used in the frozen repricing configurations supporting Table 7.9. The second factor is active in most intervals but becomes small in the 1–2Y bucket and hits the lower λ bound in the 5–7Y bucket, while the displacement and volatility-of-volatility terms vary substantially from bucket to bucket. These near-boundary values are part of the effective-dimensionality evidence behind the factor-study caveats in Section 7.3.8.

Table B2: Piecewise d=2 full QG-SV parameters from the 2026-03-17 calibration diagnostics.

Bucket	**End (Y)**	**ATM fwd (%)**	λ_1	λ_2	b_1	b_2	η
0–1	1.000	3.036	0.366	0.558	0.528	0.021	2.032
1–2	2.003	3.123	0.253	0.006	0.366	0.465	2.366
2–3	3.003	3.206	0.342	0.332	0.396	0.484	2.619
3–5	5.003	3.367	0.526	0.676	0.369	0.298	1.865
5–7	7.005	3.445	0.363	0.001	0.300	0.514	1.570
7–10	10.008	3.469	0.614	0.180	0.299	0.680	1.908

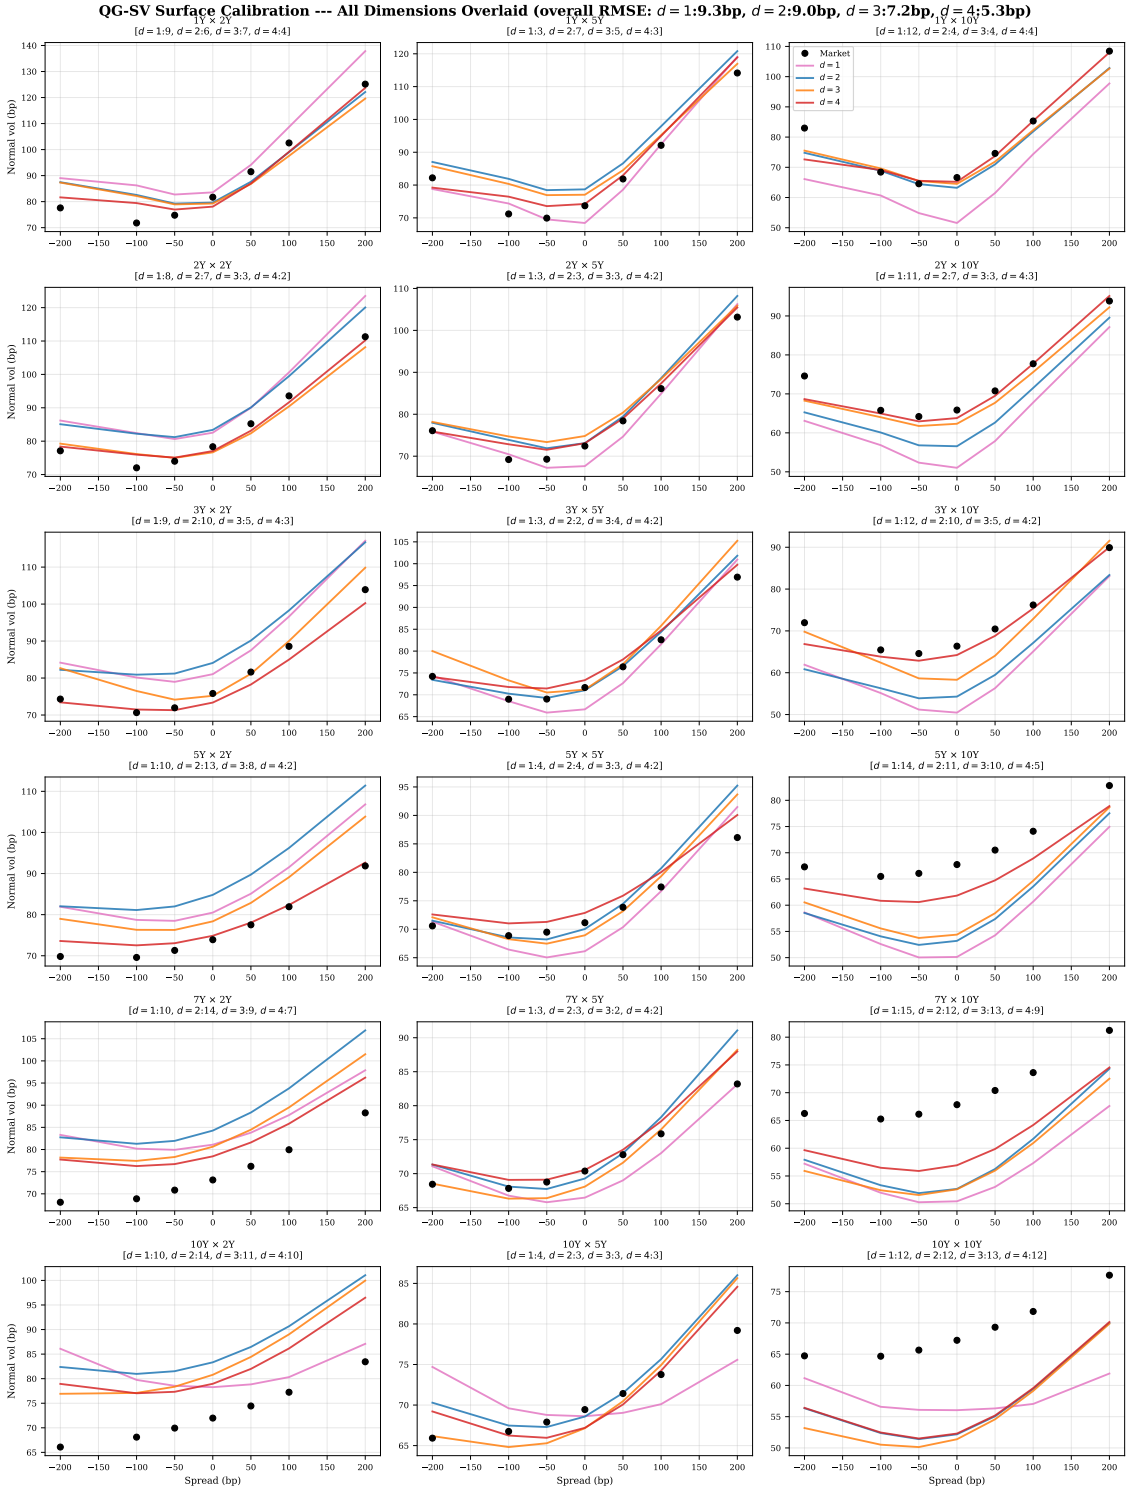

Figure B1: Full multi-tenor surface calibration with $d = 1, \ldots, 4$ overlaid on the same 2026-03-17 market data (black circles) for every expiry-tenor combination. Per-panel RMSE (in bp) is shown for each dimension in the panel title. The $d = 4$ curves cluster most tightly around the market in most cells and give the lowest reported aggregate RMSE over the tested range.

B.4 Strike-Level Smile Data for the Full QG-SV Run

Tables B3 and B4 list the 42 calibration points for the d=2 full QG-SV run. Volatilities and errors are reported in basis points of normal implied volatility. The full model fits most strikes to within a few basis points, with the largest residuals in the far upper wings of the 5Y, 7Y, and 10Y smiles.

Table B3: Saved strike-level smile data for the d=2 full QG-SV run: early expiries (1Y, 2Y, 3Y).

Expiry	$K - K_{\text{ATM}}$ **(bp)**	**Market vol (bp)**	**Model vol (bp)**	**Error (bp)**
1Y	-200	83.0	83.3	0.3
	-100	68.4	67.8	-0.6
	-50	64.5	65.5	0.9
	+0	66.6	66.6	-0.0
	+50	74.6	74.7	0.1
	+100	85.3	85.8	0.5
	+200	108.4	107.8	-0.7
2Y	-200	74.6	76.1	1.5
	-100	65.8	67.5	1.7
	-50	64.2	63.7	-0.5
	+0	65.8	64.2	-1.6
	+50	70.8	70.0	-0.8
	+100	77.7	78.4	0.7
	+200	93.8	96.8	3.0
3Y	-200	72.0	74.2	2.2
	-100	65.5	66.6	1.2
	-50	64.6	63.7	-0.9
	+0	66.3	64.5	-1.9
	+50	70.5	69.5	-1.0
	+100	76.2	76.0	-0.2
	+200	89.9	93.0	3.1

Table B4: Saved strike-level smile data for the d=2 full QG-SV run: later expiries (5Y, 7Y, 10Y).

Expiry	$K - K_{\mathrm{ATM}}$ **(bp)**	**Market vol (bp)**	**Model vol (bp)**	**Error (bp)**
5Y	-200	67.3	71.5	4.2
	-100	65.5	65.0	-0.5
	-50	66.1	63.7	-2.3
	+0	67.8	65.1	-2.6
	+50	70.5	69.5	-1.0
	+100	74.1	75.3	1.2
	+200	82.8	89.0	6.2
7Y	-200	66.3	69.5	3.3
	-100	65.3	65.2	-0.1
	-50	66.1	64.7	-1.4
	+0	67.8	66.0	-1.8
	+50	70.4	69.1	-1.3
	+100	73.6	74.0	0.3
	+200	81.2	86.2	5.0
10Y	-200	64.7	67.0	2.3
	-100	64.7	64.5	-0.2
	-50	65.6	64.4	-1.2
	+0	67.2	65.8	-1.4
	+50	69.3	68.6	-0.7
	+100	71.8	72.2	0.4
	+200	77.6	81.4	3.7