Pricing FX Quanto Range Accruals: A Comparative Study of Pricing Models

Master's thesis in Complex Adaptive Systems

EDWARD FRITZELL

DEPARTMENT OF MATHEMATICAL SCIENCES

CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

Master's Thesis MVEX03

Pricing FX Quanto Range Accruals: A Comparative Study of Pricing Models

EDWARD FRITZELL

CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Mathematical Sciences
Division of Mathematics
Chalmers University of Technology
Gothenburg, Sweden 2026

Pricing FX Quanto Range Accruals: A Comparative Study of Pricing Models
EDWARD FRITZELL

Examiner: Johan Tykesson, Department of Mathematical sciences

Master's Thesis MVEX03
Department of Mathematical Sciences
Division of Mathematics
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Typeset in LaTeX
Gothenburg, Sweden 2026

Pricing FX Quanto Range Accruals: A Comparative Study of Pricing Models
EDWARD FRITZELL
Department of Mathematical Sciences
Chalmers University of Technology

Abstract

This thesis investigates the impact of volatility modelling on the valuation of a EUR/USD range accrual quantoed to GBP. The product accrues coupon at monthly fixing dates whenever the EUR/USD spot lies inside a specified corridor, while the payoff is converted to GBP at maturity using a pre-specified quanto conversion. Since the payoff currency differs from the natural USD domestic currency of the EUR/USD underlying, valuation is performed under the GBP pricing measure. Three models are compared within a common market-data framework: a quanto-adjusted at-the-money Black–Scholes benchmark, a triangle-based Dupire local-volatility model, and the multi-factor FX Heston model of De Col, Gnoatto, and Grasselli. The local-volatility and Heston models are calibrated to EUR/USD, GBP/USD, and EUR/GBP implied-volatility smiles, and their prices are computed by Monte Carlo simulation under the GBP measure. The empirical analysis considers corridor-width and corridor-centre sweeps for maturities of 6M, 1Y, and 3Y. The results show that the volatility smile has a material and corridor-dependent effect on price. Narrow corridors are particularly sensitive, and the size and sign of the pricing difference relative to the Black–Scholes benchmark depend on both the corridor width and the corridor location relative to the spot. The local-volatility and multi-factor Heston models produce broadly similar prices, with differences that are smaller than their common deviation from the benchmark. The main conclusion is that incorporating the volatility smile is important for pricing FX quanto range accruals, while the choice between the two smile-consistent models is of secondary importance for the price level in this setting.

Keywords: FX options, quanto adjustment, range accruals, local volatility, Heston model, volatility smile, Monte Carlo.

Acknowledgements

I would like to thank my supervisor, Dmytro Fedorets, and my team at ING for welcoming me so warmly and for sharing their experience with me. Their guidance, openness, and support have been invaluable, both in helping me understand the financial industry and in giving this thesis a practical context.

Contents

1

Introduction

Derivative pricing is fundamentally about modelling uncertainty. For a foreign-exchange derivative, the price depends not only on today's exchange rate and interest-rate curves, but also on the assumed future distribution of the underlying currency pair. The classical Black–Scholes model provides a tractable lognormal benchmark for FX pricing, but observed option markets show that implied volatility varies across strikes and maturities. This volatility smile indicates that the market assigns probabilities differently from the simple lognormal model [1, 2].

This thesis studies the impact of volatility modelling on a EUR/USD range accrual quantoed to GBP. The product accrues coupon at monthly fixing dates whenever EUR/USD lies inside a specified corridor, and the payoff is converted to GBP at maturity using a pre-specified quanto conversion. Since the payoff is quantoed to GBP rather than paid in its natural USD currency, valuation must be performed under the GBP pricing measure. The quanto feature therefore changes the drift of the EUR/USD dynamics and links the pricing problem to the joint behaviour of EUR/USD and GBP/USD.

Three models are compared within the same market-data framework. The first is the Black–Scholes benchmark, which provides a transparent analytical reference but ignores the volatility smile. The second is a Dupire local-volatility model, where local-volatility surfaces are extracted from fitted implied-volatility smiles [3, 4]. The third is the multi-factor FX Heston model of De Col, Gnoatto, and Grasselli [5], which introduces stochastic volatility while preserving the triangular structure of FX markets.

The comparison is performed across corridor-width and corridor-centre sweeps for maturities of 6M, 1Y, and 3Y. This makes it possible to study not only whether the smile-consistent models differ from the benchmark, but also where the differences are largest. The results show that the volatility smile has a material and corridor-dependent effect on the range-accrual price. The largest deviations from the benchmark occur for narrow corridors, while the difference between the local-volatility and multi-factor Heston models is generally smaller than the difference between either of them and Black–Scholes.

The study is limited to one market data set, one currency triangle, and three maturities, so the results should be interpreted in that context. The remainder of the thesis is organised as follows. The Theory chapter introduces the pricing framework, the

quanto range-accrual payoff, implied volatility, and the three models. The Method chapter describes the market-data processing and numerical implementation. The Results chapter summarizes the model calibration and presents the corridor sensitivity analysis. The thesis ends with a conclusion and suggestions for further work.

2

Theory

This chapter develops the theoretical background needed for the rest of the thesis. It begins with the probabilistic and stochastic-calculus tools that underpin derivative pricing, then specialises these ideas to the no-arbitrage pricing framework and the foreign-exchange market. With this in place, the chapter introduces the quanto setting, the range-accrual payoff, the volatility smile, and finally the three models compared in the numerical work: a Garman–Kohlhagen benchmark, a local-volatility model on the foreign-exchange triangle, and a multi-factor Heston model. The chapter ends with the Monte Carlo techniques used in the Monte Carlo simulations.

2.1 Mathematical preliminaries

This section recalls the basic stochastic-calculus tools used throughout the thesis. The presentation follows [6].

2.1.1 Filtered probability spaces and Brownian motion

The starting point is a filtered probability space

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

where the filtration $\{\mathcal{F}_t\}_{t\geq 0}$ represents the information available up to time t. A stochastic process $X = \{X_t\}_{t\geq 0}$ is adapted if X_t is $\mathcal{F}_t$-measurable for every t [6]. In a financial context, adaptedness expresses the natural property that the value of a process at time t depends only on information available at that time.

A standard Brownian motion $W = \{W_t\}_{t\geq 0}$ is a continuous adapted Gaussian process with $W_0 = 0$, independent increments, and

$$W_t - W_s \sim N(0, t-s), \qquad 0 \leq s < t.$$

Brownian motion serves as the basic source of randomness in all models considered in this thesis.

2.1.2 Stochastic differential equations and Itô calculus

The models used in this thesis describe market quantities as stochastic processes evolving continuously in time. The natural mathematical language for such models

is stochastic differential equations [6, 7]. A one-dimensional diffusion has the form

$$\mathrm{d}X_t = \mu(X_t, t)\,\mathrm{d}t + \sigma(X_t, t)\,\mathrm{d}W_t, \tag{2.1}$$

or equivalently

$$X_t = X_0 + \int_0^t \mu(X_s, s)\,\mathrm{d}s + \int_0^t \sigma(X_s, s)\,\mathrm{d}W_s.$$

The drift μ describes the predictable part of the motion, while the diffusion coefficient σ determines the size of the random fluctuations. In derivative pricing, the drift is determined by the pricing measure, while the diffusion term represents the chosen volatility model.

Since stochastic processes driven by Brownian motion are not differentiable in the ordinary sense, the usual chain rule cannot be applied directly. Instead, transformations of stochastic processes are handled using Itô's formula, stated next.

Theorem 2.1.1 (Itô's formula) *Let X_t satisfy (2.1) and let $f \in C^{1,2}([0,T] \times \mathbb{R}; \mathbb{R})$. Then*

$$\mathrm{d}f(X_t, t) = \Big(\frac{\partial f}{\partial t} + \mu(X_t, t)\frac{\partial f}{\partial x} + \frac{1}{2}\sigma(X_t, t)^2\frac{\partial^2 f}{\partial x^2}\Big)\mathrm{d}t + \sigma(X_t, t)\frac{\partial f}{\partial x}\,\mathrm{d}W_t.$$

The proof is found in [6, 7]. The additional second-order term is the main difference between Itô's formula and the ordinary chain rule.

A useful special case for derivative pricing is the logarithm of the diffusion. If S_t follows the multiplicative dynamics

$$\mathrm{d}S_t = \mu_t S_t\,\mathrm{d}t + \sigma_t S_t\,\mathrm{d}W_t,$$

then applying Theorem 2.1.1 with $f(S) = \ln S$ gives

$$\mathrm{d}(\ln S_t) = \Big(\mu_t - \frac{1}{2}\sigma_t^2\Big)\mathrm{d}t + \sigma_t\,\mathrm{d}W_t.$$

This identity is used repeatedly in the pricing models below. It explains why the log-spot drift contains the correction term $-\frac{1}{2}\sigma_t^2$. It is also the basis for the log-Euler discretisations used in the Monte Carlo implementation.

2.1.3 Change of measure

Derivative pricing often requires switching from one probability measure to another. In a diffusion setting, Girsanov's theorem describes how this change affects the dynamics of an Itô process. The key consequence is that a change of equivalent measure modifies the drift but leaves the diffusion coefficient unchanged.

Theorem 2.1.2 (Girsanov) *Let $(X_t)_{t\in[0,T]}$ satisfy*

$$\mathrm{d}X_t = \mu_1(X_t, t)\,\mathrm{d}t + \sigma(X_t, t)\,\mathrm{d}B_t$$

under a measure P, where B_t is a P-Brownian motion and $\sigma(X_t, t) > 0$. Define

$$H_t = \frac{\mu_1(X_t, t) - \mu_2(X_t, t)}{\sigma(X_t, t)},$$

and the equivalent measure Q with Radon–Nikodym density

$$\frac{\mathrm{d}Q}{\mathrm{d}P} = \exp\Big(-\int_0^T H_s \,\mathrm{d}B_s - \frac{1}{2}\int_0^T H_s^2 \,\mathrm{d}s\Big).$$

Then

$$W_t = B_t + \int_0^t H_s \,\mathrm{d}s$$

is a Brownian motion under Q, and X_t satisfies

$$\mathrm{d}X_t = \mu_2(X_t, t)\,\mathrm{d}t + \sigma(X_t, t)\,\mathrm{d}W_t.$$

The proof is found in [6]. Theorem 2.1.2 motivates an important rule of thumb in derivative pricing. When valuing a payoff under a different numéraire or pricing currency, only the drift of the underlying needs to be adjusted. The diffusion specification carries over unchanged. This observation is central to the foreign-exchange and quanto pricing sections below.

2.2 No-arbitrage pricing

Derivative pricing rests on the principle of no arbitrage. The fundamental theorem of asset pricing states that an arbitrage-free market admits an equivalent probability measure under which discounted traded assets are martingales [7]. The associated measure is called the risk-neutral measure or pricing measure, and the choice of measure is tied to the choice of numéraire.

Under the pricing measure $\mathbb{Q}$ associated with a chosen pricing currency, the time-zero value of a payoff $V(T)$ is

$$V(0) = \mathbb{E}^{\mathbb{Q}}\Big[\exp\Big(-\int_0^T r(u)\,\mathrm{d}u\Big)V(T)\Big],$$

where $r(t)$ is the instantaneous short rate of the pricing currency. When the short rate is deterministic, the discount factor can be pulled out of the expectation,

$$V(0) = P(0,T)\,\mathbb{E}^{\mathbb{Q}}\big[V(T)\big], \tag{2.2}$$

where

$$P(0,T) = \exp\Big(-\int_0^T r(u)\,\mathrm{d}u\Big)$$

is the time-zero price of a zero-coupon bond paying one unit of the pricing currency at time T.

The pricing rule (2.2) reduces derivative valuation to two tasks, discounting and evaluating an expectation under the appropriate pricing measure. By the fundamental theorem, the drift of the underlying under this measure is fixed by no arbitrage. Differences between models therefore enter through the volatility specification and the corresponding model-specific quanto drift, a point that becomes the focus of Section 2.7.

2.3 Foreign-exchange dynamics

This section specialises the pricing framework of Section 2.2 to foreign-exchange markets and introduces the triangle relations that constrain volatility models on more than one currency pair.

Foreign-exchange rates are quoted as the price of one currency in terms of another. In this thesis,

$$S_t = \text{EUR/USD}$$

denotes the number of US dollars paid for one euro at time t. This is the spot exchange rate. A higher value of S_t therefore means that the euro has strengthened relative to the dollar [2].

The forward exchange rate $F(0,T)$ is the rate agreed at time zero for an exchange of currencies at time T. Under no arbitrage, the forward rate is determined by the spot rate and the discount factors of the two currencies [2]. For a pair with domestic currency d and foreign currency f,

$$F(0,T) = S_0 \frac{P_f(0,T)}{P_d(0,T)}.$$

For the EUR/USD pair, USD is the domestic currency and EUR is the foreign currency, so

$$F(0,T) = S_0 \frac{P_{\text{EUR}}(0,T)}{P_{\text{USD}}(0,T)}. \tag{2.3}$$

Under the USD risk-neutral measure, the spot dynamics of EUR/USD follow

$$\frac{\mathrm{d}S_t}{S_t} = \Big(r_{\text{USD}}(t) - r_{\text{EUR}}(t)\Big)\,\mathrm{d}t + \sigma_t\,\mathrm{d}W_t^{\text{USD}},$$

where r_{USD} and r_{EUR} are the USD and EUR short rates [1, 2]. The drift equals the difference of the two short rates and ensures consistency with the FX forward formula (2.3).

A central feature of FX markets is that exchange rates are linked through currency triangles. For three currencies such as EUR, USD, and GBP, the spot rates satisfy

$$\text{EUR/GBP} = \frac{\text{EUR/USD}}{\text{GBP/USD}}. \tag{2.4}$$

The same identity must hold for forward rates under no arbitrage. A model that violates this relation would generate risk-free profit by trading around the triangle. Triangle consistency is therefore not only a property of spot and forward rates, but also a guiding principle for volatility modelling across pairs, an idea developed in Section 2.7.

2.4 Quanto pricing framework

A quanto option is a derivative whose payoff is converted into a currency different from the natural payoff currency of the underlying using a pre-specified conversion rate [8, 2]. This conversion rate is fixed at the trade date and may, for example, be normalised to one. In the numerical work of this thesis, the underlying is EUR/USD, whose natural domestic currency is USD, while the range-accrual payoff is quantoed to GBP.

By the change-of-numéraire principle, valuation must be performed under the pricing measure of the payoff currency. Hence, for a payoff $\Pi(T)$ after conversion into GBP at maturity T, the time-zero value is

$$V(0) = P_{\mathrm{GBP}}(0,T)\mathbb{E}^{\mathbb{Q}^{\mathrm{GBP}}}\big[\Pi(T)\big], \tag{2.5}$$

where $P_{\mathrm{GBP}}(0,T)$ is the GBP discount factor and $\mathbb{Q}^{\mathrm{GBP}}$ is the corresponding pricing measure.

The choice of measure has two consequences for the FX underlying. First, all probabilities and expectations involving S_t must be evaluated under $\mathbb{Q}^{\mathrm{GBP}}$, not under its natural USD measure. Second, by Theorem 2.1.2 the dynamics of S_t under $\mathbb{Q}^{\mathrm{GBP}}$ differ from those under the USD measure through a change of drift, while the diffusion specification carries over unchanged. The quanto feature is therefore not a correction applied after pricing, but a change of pricing measure that affects the dynamics of the underlying.

The explicit form of the EUR/USD drift under $\mathbb{Q}^{\mathrm{GBP}}$ depends on the model. In the Black–Scholes benchmark it reduces to a deterministic adjustment involving a single correlation, while in the local-volatility and multi-factor Heston models the drift is derived inside a higher-dimensional FX framework. The model-specific forms appear in Section 2.7.

2.5 Range accruals

A range accrual is a structured product that pays a coupon proportional to the fraction of fixing dates on which the underlying remains within a prescribed corridor [2]. It is the product priced in the numerical work of this thesis, in the GBP-quantoed form introduced in Section 2.4.

Let $t_1 < \cdots < t_N$ be the fixing dates, K_{low} and K_{up} the lower and upper barriers, and c the coupon rate. After the quanto conversion, with the conversion factor normalised to one, the payoff at maturity is

$$\Pi(T) = c \cdot \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\{K_{\mathrm{low}} < S_{t_i} < K_{\mathrm{up}}\}. \tag{2.6}$$

At each fixing date a unit accrues if the underlying lies inside the corridor, and the average is paid at maturity.

Substituting (2.6) with $c = 1$ into the quanto pricing formula (2.5) and using the linearity of expectation gives

$$\mathrm{PV} = P_{\mathrm{GBP}}(0,T) \cdot \frac{1}{N} \sum_{i=1}^{N} \mathbb{E}^{\mathbb{Q}^{\mathrm{GBP}}}\Big[\mathbf{1}\{K_{\mathrm{low}} < S_{t_i} < K_{\mathrm{up}}\}\Big].$$

The expectation of an indicator equals the probability of the corresponding event, so the present value can be written as

$$\mathrm{PV} = P_{\mathrm{GBP}}(0,T) \cdot \frac{1}{N} \sum_{i=1}^{N} \mathbb{Q}^{\mathrm{GBP}}\Big(K_{\mathrm{low}} < S_{t_i} < K_{\mathrm{up}}\Big). \tag{2.7}$$

The corridor event itself can be split into the difference of two one-sided events,

$$\mathbf{1}\{K_{\mathrm{low}} < S_t < K_{\mathrm{up}}\} = \mathbf{1}\{S_t > K_{\mathrm{low}}\} - \mathbf{1}\{S_t > K_{\mathrm{up}}\},$$

which gives

$$\mathbb{Q}^{\mathrm{GBP}}\Big(K_{\mathrm{low}} < S_{t_i} < K_{\mathrm{up}}\Big) = \mathbb{Q}^{\mathrm{GBP}}\Big(S_{t_i} > K_{\mathrm{low}}\Big) - \mathbb{Q}^{\mathrm{GBP}}\Big(S_{t_i} > K_{\mathrm{up}}\Big).$$

Each fixing therefore reduces to two probabilities under $\mathbb{Q}^{\mathrm{GBP}}$ [2]. The price of the range accrual is therefore determined by the marginal corridor probabilities at the fixing dates, evaluated under the GBP pricing measure.

2.6 Volatility smile and implied volatility

The models compared in the numerical work differ mainly in how they specify volatility. To make this comparison precise, this section introduces implied volatility, the implied-volatility surface, and the parametrisation used later in the local-volatility framework. The presentation does not require prior familiarity with the Black–Scholes model. The full Black–Scholes framework is introduced as one of the three pricing models in Section 2.7.1.

2.6.1 European options and implied volatility

In foreign-exchange markets, the most actively quoted instruments after the spot itself are European call and put options [2]. A European call on EUR/USD with strike K and maturity T pays $\max(S_T - K, 0)$ at T, and a European put pays $\max(K - S_T, 0)$. Their market prices reflect how participants value the future distribution of the exchange rate.

An option is called at-the-money, or ATM, when its strike is close to the current forward level of the underlying for the relevant maturity. In the FX market, the precise ATM strike depends on the quotation convention, but conceptually the ATM quote represents the centre of the volatility smile. Quotes such as 25C and 25P then describe volatility away from this centre, on the call and put sides of the smile.

To translate option prices in dollars into a single comparable number, the market uses the Garman–Kohlhagen pricing formula, the foreign-exchange version of the Black–Scholes formula [1, 2]. The formula appears in full in Section 2.7.1. For the present discussion only its structure matters. Given the spot S_0, the strike K, the maturity T, the two short rates r_d and r_f, and a volatility σ, the formula returns a model price for the option,

$$\text{price} = C_{\text{GK}}\Big(S_0, K, T, r_d, r_f, \sigma\Big). \tag{2.8}$$

All inputs except the volatility are directly observable in the market. A market price therefore pins down a unique value of σ through (2.8).

The implied volatility $\sigma_{\text{imp}}(K, T)$ is defined as the value of σ that makes the Garman–Kohlhagen price match the observed market price. For a given strike and maturity, it solves

$$C_{\text{mkt}}(K, T) = C_{\text{GK}}\Big(S_0, K, T, r_d, r_f, \sigma_{\text{imp}}(K, T)\Big).$$

Implied volatility is therefore not directly observable but is obtained by inverting the pricing formula numerically.

The convenience of this convention is that every option price in the market can be reported as a single dimensionless number rather than as a currency amount that depends on spot, strike, and rates. If the market priced all options consistently with the assumptions of the Garman–Kohlhagen model, in particular that the spot is lognormal with constant volatility, then $\sigma_{\text{imp}}(K, T)$ would be the same number for every K and T.

Empirically this is not the case. Implied volatility varies systematically with strike, generating the implied-volatility smile, and with maturity, generating the implied-volatility term structure. Collecting these values across strikes and maturities defines the implied-volatility surface $\sigma_{\text{imp}}(K, T)$ [4, 9]. The shape of the surface is the empirical fingerprint of how the market deviates from the lognormal assumption, and reproducing it is the central modelling challenge addressed by the local-volatility and Heston models in Section 2.7.

In FX markets, implied volatilities are often quoted by option delta rather than directly by strike. The delta measures the sensitivity of the option price to the spot exchange rate and provides a moneyness convention commonly used by market participants. In the implementation, these delta quotes are converted to strikes before constructing the implied-volatility surface.

2.6.2 Log-forward moneyness and total variance

Following [4], the strike dimension of the surface is conveniently expressed in terms of log-forward moneyness,

$$y = \ln\Big(\frac{K}{F(T)}\Big),$$

where $F(T)$ is the forward price for maturity T. This transformation centres the smile around the forward and yields a natural coordinate system for comparing smiles across maturities.

It is also convenient to work with total implied variance,

$$w(y,T) = \sigma_{\text{imp}}(K,T)^2\,T,$$

rather than with implied volatility itself. Total variance is the natural quantity in local-volatility modelling, since Dupire's formula in the Gatheral form, given later in Section 2.7.2, is expressed directly in terms of derivatives of $w(y,T)$. Constructing a smooth and arbitrage-consistent surface in $w(y,T)$ is therefore a central step in the local-volatility framework.

2.7 Pricing models

This section presents the three models compared in the numerical work. They share the same FX underlying and the same GBP pricing framework introduced in Section 2.4, but they differ in how the volatility is specified, how the smile is captured, and how the quanto adjustment enters the dynamics. The Black–Scholes benchmark uses only at-the-money information. The local-volatility model fits the entire implied-volatility surface but keeps volatility deterministic. The multi-factor Heston model introduces stochastic variance factors in a triangle-consistent manner.

In the models, the underlying pair EUR/USD is denoted by $S_1(t)$, while $S_2(t)$ denotes GBP/USD and $S_{12}(t)$ denotes EUR/GBP.

2.7.1 The Black–Scholes benchmark

The Garman–Kohlhagen model is the foreign-exchange version of the Black–Scholes framework [1, 2]. It assumes that the exchange rate follows lognormal dynamics with deterministic, possibly time-dependent volatility. Under the USD pricing measure, the EUR/USD spot satisfies

$$\frac{\mathrm{d}S_1(t)}{S_1(t)} = \Big(r_{\text{USD}}(t) - r_{\text{EUR}}(t)\Big)\,\mathrm{d}t + \sigma_1(t)\,\mathrm{d}W_1^{\text{USD}}(t).$$

Since the payoff is quantoed to GBP, valuation must instead be performed under the GBP pricing measure. In the one-factor lognormal setting, the corresponding quanto-adjusted dynamics are

$$\frac{\mathrm{d}S_1(t)}{S_1(t)} = \Big[r_{\text{USD}}(t) - r_{\text{EUR}}(t) + \rho_{1,2}(t)\,\sigma_1(t)\,\sigma_2(t)\Big]\mathrm{d}t + \sigma_1(t)\,\mathrm{d}W_1^{\text{GBP}}(t). \tag{2.9}$$

Here $\sigma_1(t)$ is the instantaneous EUR/USD volatility, $\sigma_2(t)$ is the instantaneous GBP/USD volatility, and $\rho_{1,2}(t)$ is the instantaneous correlation between EUR/USD and GBP/USD [8, 2]. Throughout, parameters written as functions of time in the

dynamics denote instantaneous quantities, that is, the coefficients appearing in the stochastic differential equation at the instant t. Where a closed-form pricing expression requires the corresponding quantity averaged over the interval $[0, t_i]$, this is called an effective quantity and is written with a bar. The additional drift term $\rho_{1,2}(t)\,\sigma_1(t)\,\sigma_2(t)$ is the quanto adjustment generated by changing from the USD to the GBP pricing measure. The instantaneous correlation can be inferred from the instantaneous volatilities through the triangle relation

$$\rho_{1,2}(t) = \frac{\sigma_1(t)^2 + \sigma_2(t)^2 - \sigma_{12}(t)^2}{2\,\sigma_1(t)\,\sigma_2(t)}.$$

A deterministic time-dependent volatility does not change the lognormal structure of the model. The important point is that the distribution at a fixing date depends on the accumulated variance over the interval $[0, t_i]$, not on the instantaneous volatility at the fixing date [2]. Define the total EUR/USD variance by

$$w_1(t) = \int_0^t \sigma_1^2(u)\,\mathrm{d}u. \tag{2.10}$$

Equivalently, the effective Black volatility over $[0, t]$ is

$$\bar{\sigma}_1(t) = \sqrt{\frac{w_1(t)}{t}},$$

the constant volatility which produces the same total variance over $[0, t]$. By the Itô isometry [7], $w_1(t)$ is precisely the variance of $\ln S_1(t)$ under the time-dependent dynamics, so it is the quantity that enters the lognormal distribution rather than $\sigma_1(t)^2 t$.

The drift entering the lognormal distribution is also an average over the pricing interval. Define the effective GBP-measure drift by

$$\bar{\mu}_1^{\mathrm{GBP}}(t) = \frac{1}{t}\int_0^t \Big[r_{\mathrm{USD}}(u) - r_{\mathrm{EUR}}(u) + \rho_{1,2}(u)\,\sigma_1(u)\,\sigma_2(u)\Big]\,\mathrm{d}u. \tag{2.11}$$

The rate component of (2.11) is obtained directly from the discount curves. Since $P_c(0,t) = \exp\Big(-\int_0^t r_c(u)\,\mathrm{d}u\Big)$ for each currency c, the time-averaged rate differential is

$$\frac{1}{t}\int_0^t \Big[r_{\mathrm{USD}}(u) - r_{\mathrm{EUR}}(u)\Big]\,\mathrm{d}u = \frac{1}{t}\,\ln\frac{P_{\mathrm{EUR}}(0,t)}{P_{\mathrm{USD}}(0,t)}. \tag{2.12}$$

The quanto component of (2.11) is the time-average of the product $\rho_{1,2}(u)\,\sigma_1(u)\,\sigma_2(u)$. Rather than estimating the instantaneous correlation, this integrated covariance is obtained directly from the FX triangle. This gives the integrated quanto drift in terms of the three total-variance term structures,

$$\int_0^t \rho_{1,2}(u)\,\sigma_1(u)\,\sigma_2(u)\,\mathrm{d}u = \tfrac{1}{2}\Big[w_1(t) + w_2(t) - w_{12}(t)\Big], \tag{2.13}$$

so the effective quanto adjustment is

$$\bar{q}_{1,2}(t) = \frac{1}{2t}\Big[w_1(t) + w_2(t) - w_{12}(t)\Big]. \tag{2.14}$$

Within this ATM total-variance construction, equation (2.14) gives the corresponding time-average of the quanto covariance term, without requiring a separate estimate of the instantaneous correlation. It uses only the ATM total-variance term structures of the three pairs.

Applying Itô's formula to $\ln S_1(t)$ in (2.9) gives

$$\ln S_1(t) \sim \mathcal{N}\Big(\ln S_1(0) + \bar{\mu}_1^{\mathrm{GBP}}(t)\, t - \tfrac{1}{2} w_1(t),\ w_1(t)\Big)$$

under $\mathbb{Q}^{\mathrm{GBP}}$. It follows that the corridor probability at fixing date t_i is

$$\mathbb{Q}^{\mathrm{GBP}}(K_{\mathrm{low}} < S_1(t_i) < K_{\mathrm{up}}) = N(d_2(K_{\mathrm{low}}, t_i)) - N(d_2(K_{\mathrm{up}}, t_i))\,, \tag{2.15}$$

with

$$d_2(K, t_i) = \frac{\ln(S_1(0)/K) + \bar{\mu}_1^{\mathrm{GBP}}(t_i)\, t_i - \frac{1}{2} w_1(t_i)}{\sqrt{w_1(t_i)}}. \tag{2.16}$$

Every quantity in (2.16) is an effective quantity over $[0, t_i]$. The instantaneous coefficients of (2.9) enter only through the time-integrals (2.10), (2.12), and (2.13). Substituting (2.15) into the range-accrual pricing formula (2.7) gives the Black–Scholes benchmark price.

The benchmark is therefore a closed-form lognormal reference model. It captures the ATM volatility term structure through the total variance $w_1(t)$, but it does not use strike-dependent implied volatilities. Smile information is ignored by construction.

2.7.2 The local-volatility model

The Black–Scholes benchmark cannot reproduce the dependence of option prices on strike. The local-volatility framework removes this limitation by replacing the deterministic volatility with a state- and time-dependent function $\sigma_{\mathrm{loc}}(S_t, t)$, chosen so that the model exactly reproduces the observed implied-volatility surface [3].

In a one-factor FX setting under a domestic measure, the spot dynamics are

$$\frac{\mathrm{d}S_t}{S_t} = \Big(r_d(t) - r_f(t)\Big)\,\mathrm{d}t + \sigma_{\mathrm{loc}}(S_t, t)\,\mathrm{d}W_t.$$

The local-volatility function is not specified directly, but is inferred from the implied-volatility surface introduced in Section 2.6. The construction proceeds in two steps. First, a smooth and arbitrage-consistent implied-volatility surface is fitted to market data. Second, Dupire's formula is applied to extract the corresponding local volatility.

For each maturity, the implied-volatility smile is parametrised using the raw SVI form [4, 9],

$$w(y) = a + b\Big[\rho_{\mathrm{SVI}}(y - m) + \sqrt{(y - m)^2 + \sigma_{\mathrm{SVI}}^2}\Big], \tag{2.17}$$

where $w(y,T) = \sigma_{\text{imp}}(K,T)^2 T$ is the total implied variance and $(a, b, \rho_{\text{SVI}}, m, \sigma_{\text{SVI}})$ are the slice parameters. The parameters are calibrated for each tenor and FX pair, and the resulting slices are interpolated across maturity to produce a smooth surface.

Given this surface, the local volatility is recovered by Dupire's formula, stated next in the form due to Gatheral.

Theorem 2.7.1 (Dupire's formula, Gatheral form) *Let $w(y,T) = \sigma_{\text{imp}}(K,T)^2 T$ denote total implied variance as a function of log-forward moneyness $y = \ln(K/F(T))$ and maturity T. The local-volatility function consistent with the implied-volatility surface is given by*

$$\sigma^2_{\text{loc}}(K,T) = \frac{\partial_T w}{1 - \frac{y}{w}\partial_y w + \frac{1}{4}\Big(-\frac{1}{4} - \frac{1}{w} + \frac{y^2}{w^2}\Big)(\partial_y w)^2 + \frac{1}{2}\partial_{yy} w}.$$

The proof is found in [4]. The local-volatility surface is therefore determined by the fitted implied-volatility surface, so its quality depends directly on the smoothness and arbitrage consistency of that fitted surface.

To handle a quanto product on a currency pair, a single-pair local-volatility model is not sufficient. Pricing under the GBP measure couples the dynamics of EUR/USD to those of a second pair, and consistency with the FX triangle must be enforced. The model is therefore formulated on the full triangle

$$S_1 = \text{EUR/USD}, \qquad S_2 = \text{GBP/USD}, \qquad S_{12} = \text{EUR/GBP}.$$

Under the USD pricing measure, the two USD pairs follow

$$\begin{aligned}
\frac{\mathrm{d}S_1(t)}{S_1(t)} &= \Big(r_{\text{USD}}(t) - r_{\text{EUR}}(t)\Big)\,\mathrm{d}t + \sigma_1(S_1(t),t)\,\mathrm{d}W_1^{\text{USD}}(t),\\
\frac{\mathrm{d}S_2(t)}{S_2(t)} &= \Big(r_{\text{USD}}(t) - r_{\text{GBP}}(t)\Big)\,\mathrm{d}t + \sigma_2(S_2(t),t)\,\mathrm{d}W_2^{\text{USD}}(t),
\end{aligned}$$

with state-dependent instantaneous correlation

$$\mathrm{d}\langle W_1^{\text{USD}}, W_2^{\text{USD}}\rangle_t = \rho(t, S_1(t), S_2(t))\,\mathrm{d}t.$$

The choice of local-correlation candidate $\rho(t, S_1(t), S_2(t))$ is explored in [10]. In the remainder of this thesis, we define the following local-correlation candidate

$$\rho^*(t,S_1,S_2) = \frac{\sigma_1(S_1,t)^2 + \sigma_2(S_2,t)^2 - \sigma_{12}\Big(S_1/S_2, t\Big)^2}{2\,\sigma_1(S_1,t)\,\sigma_2(S_2,t)}. \tag{2.18}$$

Under the GBP pricing measure, the diffusion coefficients carry over from the USD measure by Theorem 2.1.2, while the drift of S_1 acquires a stochastic quanto adjustment. For readability, write

$$\sigma_{1,t} = \sigma_1(S_1(t),t), \qquad \sigma_{2,t} = \sigma_2(S_2(t),t), \qquad \rho_t = \rho(t, S_1(t), S_2(t)).$$

Then the GBP-measure dynamics can be written as

$$\frac{\mathrm{d}S_1(t)}{S_1(t)} = \Big(r_{\mathrm{USD}}(t) - r_{\mathrm{EUR}}(t) + \rho_t \sigma_{1,t}\sigma_{2,t}\Big)\mathrm{d}t + \sigma_{1,t}\,\mathrm{d}W_1^{\mathrm{GBP}}(t). \tag{2.19}$$

The second state variable $S_2 = \mathrm{GBP/USD}$ must also be expressed under the GBP pricing measure, since it is simulated together with S_1. Under $\mathbb{Q}^{\mathrm{GBP}}$, the dynamics of S_2 is written as

$$\frac{\mathrm{d}S_2(t)}{S_2(t)} = \Big(r_{\mathrm{USD}}(t) - r_{\mathrm{GBP}}(t) + \sigma_{2,t}^2\Big)\mathrm{d}t + \sigma_{2,t}\,\mathrm{d}W_2^{\mathrm{GBP}}(t). \tag{2.20}$$

Applying Itô's formula to the GBP-measure dynamics gives the corresponding log-spot equations

$$\begin{aligned}
\mathrm{d}\ln S_1(t) &= \Big[r_{\mathrm{USD}}(t) - r_{\mathrm{EUR}}(t) + \rho_t\sigma_{1,t}\sigma_{2,t} - \tfrac{1}{2}\sigma_{1,t}^2\Big]\mathrm{d}t + \sigma_{1,t}\,\mathrm{d}W_1^{\mathrm{GBP}}(t),\\
\mathrm{d}\ln S_2(t) &= \Big[r_{\mathrm{USD}}(t) - r_{\mathrm{GBP}}(t) + \tfrac{1}{2}\sigma_{2,t}^2\Big]\mathrm{d}t + \sigma_{2,t}\,\mathrm{d}W_2^{\mathrm{GBP}}(t).
\end{aligned}$$

The change of probability measure for the local volatility model is explained in [10], however for this thesis, it is enough to know the dynamics under $\mathbb{Q}^{\mathrm{GBP}}$ and not the derivation.

In contrast to the deterministic Black–Scholes adjustment, the local-volatility quanto term depends on the joint state (S_1, S_2) and is therefore stochastic.

2.7.3 The multi-factor Heston model

The local-volatility model reproduces the implied-volatility surface exactly, but it assumes that volatility itself is deterministic. Empirically, financial returns exhibit volatility clustering, namely periods of large fluctuations are followed by further large fluctuations, and quiet periods by further quiet periods, a stylised fact documented across asset classes including foreign exchange [11]. The Heston framework addresses this by modelling volatility as a stochastic process. In a foreign-exchange setting, the model must additionally preserve no-arbitrage relations between currency pairs.

To guarantee triangle consistency, the model of De Col, Gnoatto, and Grasselli [5] introduces exchange rates relative to an artificial reference currency, rather than modelling each observable currency pair independently. Let $S_{0,i}(t)$ denote the value of currency i in units of an artificial reference currency labelled 0. Each currency follows

$$\frac{\mathrm{d}S_{0,i}(t)}{S_{0,i}(t)} = \big(r_0 - r_i\big)\,\mathrm{d}t - \boldsymbol{a}^{i\top}\sqrt{\mathrm{Diag}(\boldsymbol{V}(t))}\,\mathrm{d}\boldsymbol{Z}(t).$$

Here r_i is the short rate of currency i, d is the number of stochastic variance factors, and

$$\boldsymbol{a}^i = (a_1^i, \ldots, a_d^i)^\top \in \mathbb{R}^d$$

is the loading vector of currency i. The vector

$$\boldsymbol{Z}(t) = (Z_1(t), \dots, Z_d(t))^\top$$

contains the Brownian motions driving the currency factors. The stochastic variance factors are collected in

$$\boldsymbol{V}(t) = (V_1(t), \dots, V_d(t))^\top,$$

and

$$\mathrm{Diag}(\boldsymbol{V}(t)) = \begin{pmatrix} V_1(t) & 0 & \cdots & 0 \\ 0 & V_2(t) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & V_d(t) \end{pmatrix}.$$

Thus $\sqrt{\mathrm{Diag}(\boldsymbol{V}(t))}$ is the diagonal matrix with entries $\sqrt{V_k(t)}$. The loading vector $\boldsymbol{a}^i$ determines how strongly currency i is exposed to each variance factor.

Each variance factor follows a square-root diffusion of Cox–Ingersoll–Ross type,

$$\mathrm{d}V_k(t) = \kappa_k\Big(\theta_k - V_k(t)\Big)\,\mathrm{d}t + \xi_k\sqrt{V_k(t)}\,\mathrm{d}W_k(t), \tag{2.21}$$

which introduces both mean reversion and stochastic fluctuations of volatility, as in the classical Heston model. The parameter κ_k is the mean-reversion speed, θ_k is the long-run variance level, and ξ_k is the volatility of variance. The Brownian motion $W_k(t)$ drives the k-th variance factor.

The Brownian motions driving the currency factors and the variance factors are correlated through

$$\mathrm{d}\langle Z_k, W_h\rangle_t = \rho_k\delta_{kh}\,\mathrm{d}t.$$

Here ρ_k is the correlation between the k-th currency Brownian motion Z_k and the k-th variance Brownian motion W_k, while δ_{kh} is the Kronecker delta. Hence the correlation is factorwise: Z_k is correlated with W_k, but not with W_h for $h \neq k$.

Observable exchange rates are obtained as ratios of these artificial-reference exchange rates,

$$S_{i,j}(t) = \frac{S_{0,j}(t)}{S_{0,i}(t)},$$

which makes the FX triangle relation hold automatically. Under the domestic pricing measure Q^i, the dynamics of the pair $S_{i,j}$ become [5]

$$\frac{\mathrm{d}S_{i,j}(t)}{S_{i,j}(t)} = \Big(r_i - r_j\Big)\,\mathrm{d}t + \Big(\boldsymbol{a}^i - \boldsymbol{a}^j\Big)^\top\sqrt{\mathrm{Diag}(\boldsymbol{V}(t))}\,\mathrm{d}\boldsymbol{Z}^{Q^i}(t).$$

This construction preserves the triangular relation between currency pairs because all exchange rates are generated from the same underlying currency processes.

For an FX underlying with USD-denominated quotation and a payoff quantoed to GBP, define the loading differences

$$\boldsymbol{b} = \boldsymbol{a}^{\mathrm{USD}} - \boldsymbol{a}^{\mathrm{EUR}}, \qquad \boldsymbol{c} = \boldsymbol{a}^{\mathrm{USD}} - \boldsymbol{a}^{\mathrm{GBP}}.$$

The vector $\boldsymbol{b}$ is the factor loading of EUR/USD under the USD domestic measure, while $\boldsymbol{c}$ is the corresponding loading of GBP/USD. Since the product payoff is quantoed to GBP, the EUR/USD dynamics must be expressed under the GBP pricing measure. Following the change of probability measure in [5], this gives

$$\frac{\mathrm{d}S_t}{S_t} = \Big(r_{\mathrm{USD}}(t) - r_{\mathrm{EUR}}(t) + \boldsymbol{b}^\top \mathrm{Diag}(\boldsymbol{V}(t))\boldsymbol{c}\Big)\mathrm{d}t + \boldsymbol{b}^\top \sqrt{\mathrm{Diag}(\boldsymbol{V}(t))}\,\mathrm{d}\boldsymbol{Z}^{\mathrm{GBP}}(t).$$

The additional drift term

$$\boldsymbol{b}^\top \mathrm{Diag}(\boldsymbol{V}(t))\boldsymbol{c}$$

is the stochastic quanto adjustment. Unlike the Black–Scholes case, where this adjustment is deterministic, here it depends on the variance factors and therefore evolves randomly over time.

Applying Theorem 2.1.1 to $\ln S_t$ yields the log-spot dynamics used in simulation,

$$\begin{aligned}\mathrm{d}\ln S_t = &\Big(r_{\mathrm{USD}}(t) - r_{\mathrm{EUR}}(t) + \boldsymbol{b}^\top \mathrm{Diag}(\boldsymbol{V}(t))\boldsymbol{c} - \frac{1}{2}\boldsymbol{b}^\top \mathrm{Diag}(\boldsymbol{V}(t))\boldsymbol{b}\Big)\mathrm{d}t \\ &+ \boldsymbol{b}^\top \sqrt{\mathrm{Diag}(\boldsymbol{V}(t))}\,\mathrm{d}\boldsymbol{Z}^{\mathrm{GBP}}(t).\end{aligned} \tag{2.22}$$

The equations above specify the EUR/USD spot dynamics under the GBP pricing measure conditional on the variance factors. For Monte Carlo simulation, the variance-factor dynamics must also be expressed under the same pricing measure. In the DGG framework, this changes the CIR drift parameters while leaving the diffusion coefficient unchanged. Under Q^i, the factor V_k satisfies the CIR equation (2.21) with parameters $(\kappa_k^i, \theta_k^i, \xi_k)$. Under Q^j, the diffusion coefficient ξ_k is unchanged by Theorem 2.1.2, while the mean-reversion speed and long-run level transform according to [5]

$$\kappa_k^j = \kappa_k^i + \rho_k\,\xi_k\,(a_k^j - a_k^i), \qquad \theta_k^j = \theta_k^i\,\frac{\kappa_k^i}{\kappa_k^j}, \tag{2.23}$$

where a_k^i and a_k^j denote the k-th components of the loading vectors $\boldsymbol{a}^i$ and $\boldsymbol{a}^j$.

The model parameters are the variance-factor parameters $(\kappa_k, \theta_k, \xi_k)$, the initial values $V_k(0)$, the currency loading vectors $\boldsymbol{a}^i$, and the factor correlations ρ_k. These parameters are not fixed a priori, but are calibrated to market option data. In the empirical implementation, the loading vectors, initial variance factors, variance dynamics, and factor correlations are jointly fitted to the EUR/USD, GBP/USD, and EUR/GBP implied-volatility smiles.

2.8 Monte Carlo methods

The local-volatility and multi-factor Heston pricers used in the numerical work rely on Monte Carlo simulation. This section recalls the Monte Carlo estimator and the discretisation schemes used for the spot and variance processes.

2.8.1 Monte Carlo estimator

Monte Carlo pricing approximates a risk-neutral expectation by an empirical average over simulated paths [12]. If $X^{(j)}$ denotes the discounted payoff on path j, the estimator based on M paths is

$$\widehat{V}_M = \frac{1}{M}\sum_{j=1}^{M} X^{(j)}. \tag{2.24}$$

By the law of large numbers, $\widehat{V}_M \to V$ as $M \to \infty$. Its statistical uncertainty is measured by the standard error

$$\mathrm{SE}\big(\widehat{V}_M\big) = \frac{s_M}{\sqrt{M}},$$

where s_M is the sample standard deviation of the discounted payoffs.

For the range-accrual product introduced in Section 2.5, the discounted payoff on path j takes the form

$$X^{(j)} = P_{\mathrm{GBP}}(0,T)\cdot\frac{1}{N}\sum_{i=1}^{N}\mathbf{1}\{K_{\mathrm{low}} < S_{t_i}^{(j)} < K_{\mathrm{up}}\}. \tag{2.25}$$

2.8.2 Euler–Maruyama discretisation and full truncation

For a generic diffusion

$$\mathrm{d}X_t = a(X_t,t)\,\mathrm{d}t + b(X_t,t)\,\mathrm{d}W_t,$$

the Euler–Maruyama scheme on a uniform grid $t_n = n\Delta t$ reads

$$X_{n+1} = X_n + a(X_n,t_n)\Delta t + b(X_n,t_n)\sqrt{\Delta t}\,Z_n, \qquad Z_n \sim N(0,1). \tag{2.26}$$

The drift term is scaled by Δt, while the diffusion is scaled by $\sqrt{\Delta t}$ [12].

For spot processes, the implementation applies the Euler–Maruyama idea to the log-spot dynamics derived from Itô's formula. Simulating the log-spot and then exponentiating preserves positivity of the spot by construction. State-dependent quantities are frozen over the interval $[t_n, t_{n+1}]$ and evaluated at the left endpoint t_n, in line with the Euler–Maruyama scheme in (2.26).

The variance factors in the multi-factor Heston model require additional care because of the square-root diffusion. A naive Euler scheme can produce negative variance values and introduce discretisation bias. The implementation therefore uses full-truncation Euler, where only the positive part of the variance enters the drift and diffusion coefficients [13]. The model-specific update is given in Section 3.5.

3
Method

This chapter describes how the three models from Section 2.7 are implemented numerically and applied to the same range-accrual product. It begins with the product specification and the market data shared across all three models. It then describes each model in turn: the closed-form Black–Scholes benchmark, the local-volatility Monte Carlo simulation on the FX triangle, and the multi-factor Heston Monte Carlo simulation. The three models share the discount-curve construction, the FX volatility-quote conventions, and the underlying currency triangle, but differ in how the volatility surface is processed and how the EUR/USD dynamics under $\mathbb{Q}^{\mathrm{GBP}}$ are simulated.

3.1 Product specification

The product priced in the empirical comparison is the EUR/USD range accrual quantoed to GBP introduced in Section 2.5. Rather than pricing only one fixed corridor, the empirical study prices a family of corridors in order to examine how the model differences depend on the payoff geometry. Each corridor is parameterised by its centre C and half-width h,

$$K_{\mathrm{low}} = C - h, \qquad K_{\mathrm{up}} = C + h.$$

The coupon rate is fixed at

$$c = 1.$$

The fixing schedule is monthly, so an N-month maturity has N fixing dates $t_1, \ldots, t_N$ at one-month intervals. The pricing comparison is carried out for the three maturities 6M, 1Y, and 3Y. All prices are reported in GBP, and the pricing measure is therefore $\mathbb{Q}^{\mathrm{GBP}}$ for all three models, in line with the development of Section 2.4.

Two corridor sweeps are used in the empirical comparison. In the width sweep, the corridor centre is fixed at the EUR/USD spot,

$$C = S_0 = 1.193,$$

and the half-width is varied across

$$h \in \{0.02,\ 0.035,\ 0.05,\ 0.075,\ 0.10\}.$$

This sweep studies how prices change as the corridor becomes narrower or wider while remaining centred around the spot.

In the centre sweep, the half-width is fixed at

$$h = 0.05,$$

and the corridor centre is varied over the interval

$$C \in [0.943,\ 1.543].$$

This corresponds to offsets from -0.25 to $+0.35$ relative to the spot $S_0 = 1.193$. The wider centre sweep is used to show how the absolute price behaves as the corridor is moved from far below to far above the spot. For relative price comparisons, the discussion focuses on the central part of this interval.

For each corridor specification, the same market inputs are reused. The local-volatility surfaces, the multi-factor Heston calibration, and the Black–Scholes at-the-money volatility term structures are constructed once and then held fixed across the sweep. Thus, the corridor sensitivity analysis changes only K_{low} and K_{up}, while the market data, calibration results, Monte Carlo path counts, and time grids remain unchanged.

The underlying currency triangle is

$$S_1 = \text{EUR/USD}, \qquad S_2 = \text{GBP/USD}, \qquad S_{12} = \text{EUR/GBP}.$$

The Black–Scholes model uses only S_1 directly, together with at-the-money information from S_2 and S_{12} to construct the quanto adjustment. The local-volatility and multi-factor Heston models use the full triangle. The EUR/GBP spot is taken to be the triangle-implied value $S_{12}(0) = S_1(0)/S_2(0)$ rather than a separately quoted EUR/GBP spot, so the spot triangle relation (2.4) holds exactly at time zero across all three models.

3.2 Market data and curve construction

The implementation requires three categories of market data: spot rates for EUR/USD and GBP/USD, discount curves for USD, EUR, and GBP, and FX vanilla-volatility quotes for EUR/USD, GBP/USD, and EUR/GBP. The spot rate for EUR/GBP is obtained from the spot triangle relation (2.4). The market data used is from January 29, 2026.

Discount curves are interpolated log-linearly in the discount factor. This preserves positivity of $P(0,t)$ and gives a continuous discount curve for intermediate maturities. The discount curves are used both for discounting cash flows and for constructing the deterministic rate contribution in the Monte Carlo drift. In the simulation, this rate contribution is obtained from discount-factor increments rather than by point-sampling instantaneous forward rates.

FX implied volatilities are quoted in delta space rather than directly in strike space [2]. For each currency pair and tenor, the smile is represented by five market quotes,

$$10\text{C}, \qquad 25\text{C}, \qquad \text{ATM}, \qquad 25\text{P}, \qquad 10\text{P}.$$

Here 10C and 25C denote call quotes with absolute deltas of 10% and 25%, while 25P and 10P denote put quotes with absolute deltas of 25% and 10%.

Before calibration, these delta quotes are converted to strike-volatility points using standard FX conventions described by [2]. In the implementation, EUR/USD and GBP/USD use forward-unadjusted deltas, while EUR/GBP uses forward-premium-adjusted deltas. The ATM quote is interpreted using the delta-neutral straddle convention, implemented in pips for EUR/USD and GBP/USD and in percent for EUR/GBP. These converted strike-volatility points are then used as input to the SVI fits, the Dupire extraction, and the Heston calibration.

3.3 Black–Scholes benchmark

The Black–Scholes benchmark implements the term-structure Garman–Kohlhagen model of Section 2.7.1. It provides a smile-independent reference against which the local-volatility and multi-factor Heston models are compared, using only the at-the-money volatility term structure. No structural model parameters are calibrated, and the benchmark is evaluated in closed form rather than by Monte Carlo. The construction has three steps: at-the-money quotes are converted to total-variance term structures, the effective drift and variance over each fixing interval are assembled, and the corridor probability is evaluated analytically fixing by fixing.

3.3.1 Total-variance term structures

For each pair $p \in \{1, 2, 12\}$, where $S_1 = \text{EUR/USD}$, $S_2 = \text{GBP/USD}$, and $S_{12} = \text{EUR/GBP}$, the quoted at-the-money volatility at maturity T_k is converted to total variance,

$$w_p(T_k) = \sigma_{\text{ATM},p}(T_k)^2 \, T_k.$$

The total-variance curve is interpolated linearly between quoted maturities, which is equivalent to assuming piecewise-constant instantaneous variance between neighbouring quoted tenors [2]. At each fixing date t_i, the interpolated value $w_p(t_i)$ is the accumulated variance over $[0, t_i]$, defined analogously to (2.10). The same procedure is applied independently to S_1, S_2, and S_{12}.

3.3.2 Effective drift and variance

At each fixing date t_i, the effective rate differential is read directly off the discount curves through (2.12),

$$\bar{r}_c(t_i) = -\frac{\ln P_c(0, t_i)}{t_i}, \qquad c \in \{\text{USD}, \text{EUR}, \text{GBP}\},$$

which is the time-averaged short rate over $[0, t_i]$ since the discount factor is the exponential of the integrated short rate. The effective quanto adjustment is evaluated from the FX triangle directly in total-variance space, following (2.13),

$$\bar{q}_{1,2}(t_i) = \frac{1}{2t_i}\Big[w_1(t_i) + w_2(t_i) - w_{12}(t_i)\Big].$$

Within the ATM total-variance construction, this gives the corresponding time-average of the quanto covariance term and uses only the three total-variance term structures, with no separate estimate of the instantaneous correlation. The effective GBP-measure drift over $[0, t_i]$ is then

$$\bar{\mu}_1^{\mathrm{GBP}}(t_i) = \bar{r}_{\mathrm{USD}}(t_i) - \bar{r}_{\mathrm{EUR}}(t_i) + \bar{q}_{1,2}(t_i), \tag{3.1}$$

in agreement with the effective drift (2.11).

3.3.3 Closed-form pricing

With the effective drift (3.1) and the EUR/USD total variance $w_1(t_i)$, the argument of the corridor probability is

$$d_2(K, t_i) = \frac{\ln(S_1(0)/K) + \bar{\mu}_1^{\mathrm{GBP}}(t_i)\, t_i - \frac{1}{2} w_1(t_i)}{\sqrt{w_1(t_i)}},$$

and the corridor probability at fixing date t_i follows from (2.15) as

$$p_i = N(d_2(K_{\mathrm{low}}, t_i)) - N(d_2(K_{\mathrm{up}}, t_i))\,. \tag{3.2}$$

For each corridor specification (C, h), the barriers are set to $K_{\mathrm{low}} = C - h$ and $K_{\mathrm{up}} = C + h$. The corridor probabilities (3.2) are substituted into the range-accrual price (2.7) to obtain the benchmark value.

This implementation is therefore an at-the-money, time-dependent Garman–Kohlhagen benchmark. It captures the at-the-money volatility term structure through total variance but uses no strike-dependent implied volatilities, so smile information is ignored by construction. Since the benchmark is evaluated analytically through (2.7), it carries no Monte Carlo standard error, in contrast to the local-volatility and multi-factor Heston models.

3.4 Local-volatility model

The local-volatility model implements the triangle-based local-volatility model of Section 2.7.2. It adds smile consistency on top of the Black–Scholes benchmark while keeping volatility deterministic. The purpose of the procedure is to transform quoted market smiles into simulation-ready local-volatility functions. The construction has four steps: market quotes are converted from delta to strike space, smile slices are fitted with raw SVI, the fitted implied-volatility surfaces are used to extract local-volatility surfaces through Theorem 2.7.1, and the range accrual is then priced by GBP-measure Monte Carlo with a state-dependent local correlation.

3.4.1 Smile fitting

For each pair and each retained maturity, the five smile points are converted from delta-space quotes to strike-space quotes using the pair-specific FX conventions introduced in Section 3.2. Maturities below two weeks are excluded from the SVI fit. At very short maturities, total variance is small and the Dupire extraction becomes particularly sensitive to small irregularities in the fitted smile. Since the range-accrual payoff is observed at monthly fixing dates, the one-week smile has limited direct relevance for the payoff but can introduce numerical instability in the short-end local-volatility surface.

Each retained slice is expressed in log-forward moneyness and total implied variance,

$$y = \ln\left(\frac{K}{F(T)}\right), \qquad w(y,T) = \sigma_{\text{imp}}(K,T)^2 T,$$

and is fitted with the raw SVI parametrisation (2.17).

The five parameters $(a, b, \rho_{\text{SVI}}, m, \sigma_{\text{SVI}})$ are obtained per slice and per pair. The optimisation is carried out in transformed variables that enforce positivity of the relevant SVI parameters and the standard wing restriction during calibration.

3.4.2 Surface construction

Once the slice fits are available, total variance is interpolated linearly in maturity. For $T_1 < T < T_2$,

$$w(y,T) = (1-\alpha)\, w_1(y_1) + \alpha\, w_2(y_2), \qquad \alpha = \frac{T - T_1}{T_2 - T_1},$$

with forward-adjusted moneyness on each side,

$$y_i = y + \ln\left(\frac{F(T)}{F(T_i)}\right), \qquad i = 1, 2.$$

This adjustment ensures that both neighbouring SVI slices are evaluated at the same absolute strike K, rather than at the same log-moneyness value relative to different forwards.

For maturities below the first fitted tenor, total variance is scaled linearly from the first slice. For maturities beyond the last fitted tenor, implied volatility is held flat, so total variance grows linearly in maturity.

Independent SVI slices are not guaranteed to be calendar monotone in total variance. The implementation therefore checks the fitted surface on a diagnostic moneyness grid and shifts later slices upward when necessary to restore non-decreasing total variance. This reduces the risk of negative local variances in the Dupire extraction.

3.4.3 Local-volatility extraction

Local volatility is recovered from the constructed total-variance surface by Theorem 2.7.1. The local-volatility surface is precomputed on a strike-and-maturity grid for each pair and is read off this grid by interpolation during pricing. The derivatives in log-forward moneyness and maturity are computed by finite differences.

Several safeguards are used to protect the simulation from numerical instability in the Dupire formula. If the denominator is too small, if the time derivative of total variance is non-positive, or if the resulting local variance is non-positive or non-finite, the implementation falls back to the fitted implied volatility at that point. The resulting local volatility is also bounded to a broad admissible interval. The procedure produces three local-volatility surfaces,

$$\sigma_1(S_1, t), \qquad \sigma_2(S_2, t), \qquad \sigma_{12}(S_{12}, t),$$

for EUR/USD, GBP/USD, and EUR/GBP.

3.4.4 Monte Carlo pricing

Pricing under the GBP measure follows the triangle-based dynamics (2.19) for S_1 and (2.20) for S_2. Although the payoff depends only on EUR/USD, GBP/USD must also be simulated because it enters the GBP-measure quanto drift and the triangle-implied local correlation.

At time step n, the simulation first evaluates the three local volatilities at the left endpoint t_n,

$$\sigma_{1,n} = \sigma_1(S_{1,n}, t_n), \qquad \sigma_{2,n} = \sigma_2(S_{2,n}, t_n), \qquad \sigma_{12,n} = \sigma_{12}\Big(\frac{S_{1,n}}{S_{2,n}}, t_n\Big).$$

The local-correlation candidate is then formed via the triangle relation (2.18),

$$\rho_n^* = \frac{\sigma_{1,n}^2 + \sigma_{2,n}^2 - \sigma_{12,n}^2}{2\,\sigma_{1,n}\,\sigma_{2,n}}.$$

To ensure admissible correlated shocks, the candidate is clipped pathwise to a closed sub-interval of $(-1, 1)$,

$$\widetilde{\rho}_n = \min\Big(\max(\rho_n^*, -0.999),\, 0.999\Big).$$

Correlated standard normals are produced by

$$Z_{2,n} = \widetilde{\rho}_n Z_{1,n} + \sqrt{1 - \widetilde{\rho}_n^2}\, Z_n^{\perp},$$

with $Z_{1,n}$ and $Z_n^{\perp}$ independent standard normals.

The log-spot updates are

$$\begin{aligned}\ln S_{1,n+1} &= \ln S_{1,n} + \Big[r_{\mathrm{USD}}(t_n) - r_{\mathrm{EUR}}(t_n) + \widetilde{\rho}_n\,\sigma_{1,n}\,\sigma_{2,n} - \tfrac{1}{2}\sigma_{1,n}^2\Big]\Delta t \\ &\quad + \sigma_{1,n}\sqrt{\Delta t}\,Z_{1,n}, \\ \ln S_{2,n+1} &= \ln S_{2,n} + \Big[r_{\mathrm{USD}}(t_n) - r_{\mathrm{GBP}}(t_n) + \tfrac{1}{2}\sigma_{2,n}^2\Big]\Delta t \\ &\quad + \sigma_{2,n}\sqrt{\Delta t}\,Z_{2,n}.\end{aligned}$$

For each corridor specification (C, h), monthly corridor indicators are evaluated on the EUR/USD path S_1 at the fixing dates $t_1, \ldots, t_N$, using $K_{\mathrm{low}} = C - h$ and $K_{\mathrm{up}} = C + h$. The discounted payoff on each path is the range-accrual payoff (2.25), and the price is the Monte Carlo estimator (2.24) of Section 2.8, evaluated by substituting the simulated corridor frequencies for the corridor probabilities in (2.7).

The local-volatility engine uses 200,000 Monte Carlo paths and 21 time steps per month, corresponding approximately to a daily discretisation grid with 252 time steps per year. Since the payoff is evaluated at monthly fixing dates, this gives a daily-resolution simulation between fixings.

3.5 Multi-factor Heston model

The multi-factor Heston model implements the model of De Col, Gnoatto, and Grasselli introduced in Section 2.7.3. It provides a stochastic-volatility alternative to the local-volatility engine, while preserving the joint FX triangle structure. In the empirical comparison, the model is run with $d = 2$ variance factors. With $n_{\mathrm{ccy}} = 3$ currencies, the parameter count is

$$n_{\mathrm{ccy}}d + 5d = 3 \cdot 2 + 5 \cdot 2 = 16.$$

The first term counts the entries of the three loading vectors $\boldsymbol{a}^{\mathrm{USD}}, \boldsymbol{a}^{\mathrm{EUR}}, \boldsymbol{a}^{\mathrm{GBP}} \in \mathbb{R}^2$. The second counts five scalar parameters per variance factor: the initial value $V_k(0)$, the mean-reversion speed κ_k^{USD}, the long-run level θ_k^{USD}, the volatility of variance ξ_k, and the correlation ρ_k between the variance factor and its driving Brownian motion.

3.5.1 Joint calibration

The purpose of the calibration is to choose the model parameters so that the model reproduces the observed vanilla option smiles on the three currency pairs as closely as possible. The calibration uses full smile quotes for EUR/USD, GBP/USD, and EUR/GBP at the three pricing maturities 6M, 1Y, and 3Y. Quotes are converted from delta to strike space using the same conventions as in Section 3.2.

For a given parameter vector, vanilla options are priced under the appropriate domestic measure. EUR/USD and GBP/USD options are priced under Q^{USD}, while EUR/GBP options are priced under Q^{GBP} after applying the CIR measure-change rule (2.23) to the variance-factor parameters. Vanilla prices are computed by Fourier

inversion of the model's affine characteristic function, using the pricing representation of [5]. Model prices are then converted to implied volatilities by numerically inverting the Garman–Kohlhagen formula.

The calibration minimises squared errors in implied-volatility space,

$$\Theta^\star = \arg\min_{\Theta} \sum_n \left(\sigma^{\text{model}}_{\text{imp},n}(\Theta) - \sigma^{\text{mkt}}_{\text{imp},n} \right)^2,$$

where the sum runs over all quoted strike-and-maturity combinations on the three pairs. The objective is minimised in implied-volatility space rather than in price space because implied volatilities are more directly comparable across strikes and maturities, while option prices can differ substantially in scale.

The optimisation is started from the published parameter values reported in De Col, Gnoatto, and Grasselli [5], with the JPY loading entries used as an initial proxy for GBP where no GBP entries are available. These values were calibrated on a different market and a different date, so they serve as an exogenous starting point that does not depend on manual tuning to the present data set. This reduces the risk of a hand-picked initial guess, although the non-convexity of the calibration problem means that some sensitivity to the starting point may remain. The calibration output is summarised by the overall implied-volatility root-mean-square error in basis points and by per-pair fit errors.

3.5.2 Monte Carlo pricing

After calibration, the range accrual is priced by Monte Carlo under $\mathbb{Q}^{\text{GBP}}$ following the dynamics in (2.22). With the loading differences

$$\boldsymbol{b} = \boldsymbol{a}^{\text{USD}} - \boldsymbol{a}^{\text{EUR}}, \qquad \boldsymbol{c} = \boldsymbol{a}^{\text{USD}} - \boldsymbol{a}^{\text{GBP}},$$

the EUR/USD log-spot update at time t_n is

$$\ln S_{n+1} = \ln S_n + \Big[r_{\text{USD}}(t_n) - r_{\text{EUR}}(t_n) + \sum_{k=1}^{d} b_k c_k V^+_{k,n} - \tfrac{1}{2} \sum_{k=1}^{d} b_k^2 V^+_{k,n} \Big] \Delta t + \sum_{k=1}^{d} b_k \sqrt{V^+_{k,n}} \sqrt{\Delta t}\, Z_{k,n}.$$

where $V^+_{k,n} = \max(V_{k,n}, 0)$.

The variance factors are simulated under $\mathbb{Q}^{\text{GBP}}$ using the transformed parameters $(\kappa_k^{\text{GBP}}, \theta_k^{\text{GBP}})$ obtained by applying the CIR measure-change rule (2.23) to the calibrated Q^{USD} parameters. The discretisation is the full-truncation Euler scheme,

$$V_{k,n+1} = V_{k,n} + \kappa_k^{\text{GBP}} \big(\theta_k^{\text{GBP}} - V^+_{k,n} \big) \Delta t + \xi_k \sqrt{V^+_{k,n}} \sqrt{\Delta t}\, W_{k,n}, \qquad k = 1, \dots, d.$$

The factor correlation structure is imposed through

$$W_{k,n} = \rho_k Z_{k,n} + \sqrt{1 - \rho_k^2}\, Z^\perp_{k,n},$$

with $Z_{k,n}$ and $Z^{\perp}_{k,n}$ independent standard normals. Brownian shocks are independent across factors except for the imposed factorwise correlation between $Z_{k,n}$ and $W_{k,n}$.

For each corridor specification (C, h), monthly corridor indicators are evaluated on the simulated EUR/USD path using $K_{\text{low}} = C - h$ and $K_{\text{up}} = C + h$. The discounted payoff on each path is the range-accrual payoff (2.25), and the price is the Monte Carlo estimator (2.24) of Section 2.8, evaluated by substituting the simulated corridor frequencies for the corridor probabilities in (2.7).

The multi-factor Heston engine uses 200,000 Monte Carlo paths and 21 time steps per month, corresponding approximately to a daily discretisation grid with 252 time steps per year. The same time grid is used as in the local-volatility engine, so that differences between the two Monte Carlo prices are not driven by different time discretisations.

4

Results

This chapter presents the numerical results obtained from the three pricing models introduced in Sections 3.3–3.5. The purpose is to compare how the at-the-money Black–Scholes benchmark, the triangle-based local-volatility model, and the multi-factor Heston model value the same EUR/USD range accrual quantoed to GBP.

The chapter is organised in two parts. First, the calibration quality of the two smile-consistent models is summarised. Second, the corridor sensitivity analysis is presented. The detailed calibration diagnostics can be found in Appendix A, so that the main text can focus on the pricing implications. The appendix contains the fitted implied-volatility surfaces, the extracted local-volatility surfaces, and the multi-factor Heston fit diagnostics across all calibrated tenors.

4.1 Calibration summary

Before comparing range-accrual prices, the two smile-consistent models must be checked for calibration quality. The local-volatility model is constructed from SVI fits to the EUR/USD, GBP/USD, and EUR/GBP implied-volatility smiles. The fitted SVI slices produced smooth implied-volatility surfaces across the pricing region, and the calendar and numerical safeguards described in Section 3.4 gave stable local-volatility surfaces for the Monte Carlo simulation.

The Dupire extraction is more sensitive than the implied-volatility fit itself, since it depends on derivatives of the total implied variance surface. Nevertheless, the extracted local-volatility surfaces remained well behaved on the grid used for pricing. Cases where the Dupire formula became numerically unstable were handled by the fallback procedure described in Section 3.4. The full implied-volatility and local-volatility surface plots are provided in Appendix A.

The multi-factor Heston model was calibrated jointly to the EUR/USD, GBP/USD, and EUR/GBP smiles at the maturities 6M, 1Y, and 3Y. The calibration used 45 market quotes in total, with 15 quotes per currency pair. Table 4.1 reports the resulting implied-volatility root-mean-square errors.

Currency pair	Number of quotes	RMSE (bp)
EUR/USD	15	11.29
GBP/USD	15	10.28
EUR/GBP	15	9.53
Overall	45	10.39

Table 4.1: Multi-factor Heston calibration errors in implied-volatility basis points.

The calibration errors are small for all three currency pairs. The lowest pair-level error is obtained for EUR/GBP, with an RMSE of 9.53 basis points, while the highest is obtained for EUR/USD, with an RMSE of 11.29 basis points. The overall RMSE is 10.39 basis points. This indicates that the two-factor specification is able to reproduce the main level and skew of the three implied-volatility smiles simultaneously.

The largest residual errors occur in some wing quotes, while the fit is tighter around the central part of the smiles. This is consistent with the fact that the model is calibrated jointly across three currency pairs and several maturities using a finite number of stochastic variance factors. The full Heston fit diagnostics are shown in Appendix A.

These diagnostics support the use of both smile-consistent models in the price comparison. The local-volatility model reproduces the fitted implied-volatility surfaces through the Dupire construction, while the multi-factor Heston model achieves a close joint fit with a stochastic-volatility specification. The remainder of this chapter therefore focuses on the impact of these modelling choices on the range-accrual price.

4.2 Corridor sensitivity analysis

The price comparison uses the two corridor sweeps defined in Section 3.1. First, the corridor centre is fixed at the EUR/USD spot, $S_0 = 1.193$, and the half-width is varied from 0.02 to 0.10. Second, the half-width is fixed at 0.05 and the corridor centre is varied from 0.943 to 1.543. The wide centre sweep is used for the absolute-price plots. For the relative plots, the discussion focuses on the central part of the sweep, with centres between 1.143 and 1.243, since percentage differences become less informative when the absolute price is very small.

The main question is how much the smile-consistent models depart from the at-the-money Black–Scholes benchmark. This isolates the role of the volatility smile in the pricing of the range accrual. The comparison also gives information about the difference between the two smile-consistent models, since the local-volatility and multi-factor Heston results are shown side by side in the same figures.

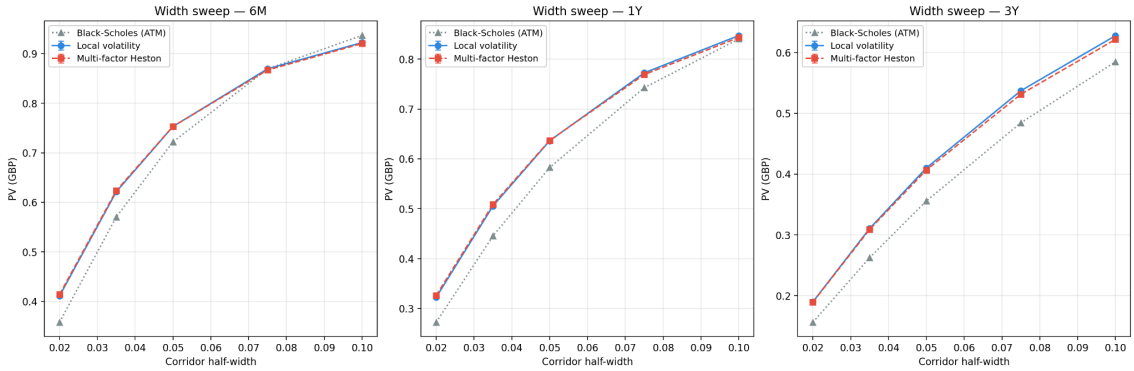

Figure 4.1: Absolute present values of the EUR/USD range accrual quantoed to GBP under the Black–Scholes benchmark, the local-volatility model, and the multi-factor Heston model, as a function of corridor half-width with the centre fixed at the EUR/USD spot. One panel is shown per maturity.

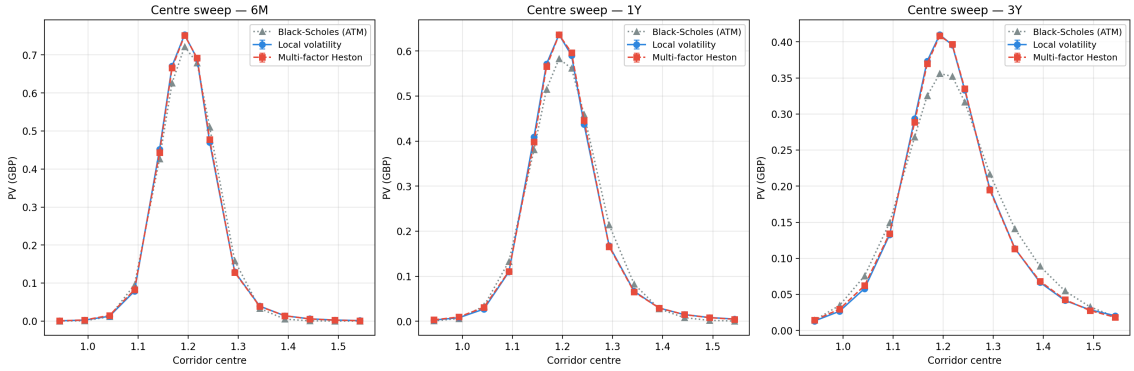

Figure 4.2: Absolute present values of the EUR/USD range accrual quantoed to GBP under the Black–Scholes benchmark, the local-volatility model, and the multi-factor Heston model, as a function of corridor centre with the half-width fixed at 0.05. One panel is shown per maturity.

Figures 4.1 and 4.2 show the absolute present values produced by all three models across the two sweeps. In the width sweep, the price increases as the corridor becomes wider. This is expected, since a wider corridor gives a higher probability that the spot lies inside the corridor at each fixing date. For a sufficiently wide corridor, this probability approaches one under all models. Therefore, since the coupon rate is normalised to one, all three prices should converge to the GBP discount factor $P_{\mathrm{GBP}}(0,T)$.

In the centre sweep, the price is highest when the corridor is centred close to the spot and falls when the corridor is moved far away. This follows directly from the payoff structure, since a corridor close to the current spot contains more probability mass than a corridor placed far in the tails. The figure also shows that the smile-consistent models produce a higher central peak than the Black–Scholes benchmark, especially for the longer maturities. This indicates that the local-volatility and Heston models assign more probability mass to the most relevant central corridor region.

The ordering between the models is not constant across the full centre sweep. In parts of the intermediate off-centre region, the Black–Scholes benchmark can produce higher prices than the smile-consistent models, while the smile-consistent prices increase relative to the benchmark again in the far-tail region. The centre sweep therefore suggests a redistribution of probability mass under the smile-consistent models: more mass near the central peak and in parts of the tails, but less in parts of the intermediate region. From a distributional perspective, the implied-volatility smile introduces asymmetry and heavier tails relative to the lognormal benchmark. This reshapes the probability density so that mass is shifted toward both the centre and the tails. Since the range-accrual payoff is effectively a sum of indicator functions, it is particularly sensitive to how probability mass is locally distributed. This explains why the smile-consistent models can produce higher prices near the spot, but lower prices in certain off-centre regions where the probability mass is reduced. Away from the central region, however, the absolute prices become small, so these tail differences are less important in present-value terms.

Figures 4.3 and 4.4 show the relative price differences of the local-volatility and multi-factor Heston models with respect to the Black–Scholes benchmark, expressed as a percentage of the benchmark price. These figures therefore show both the size of the smile correction and the similarity between the two smile-consistent models. Solid lines correspond to local volatility and dashed lines correspond to multi-factor Heston.

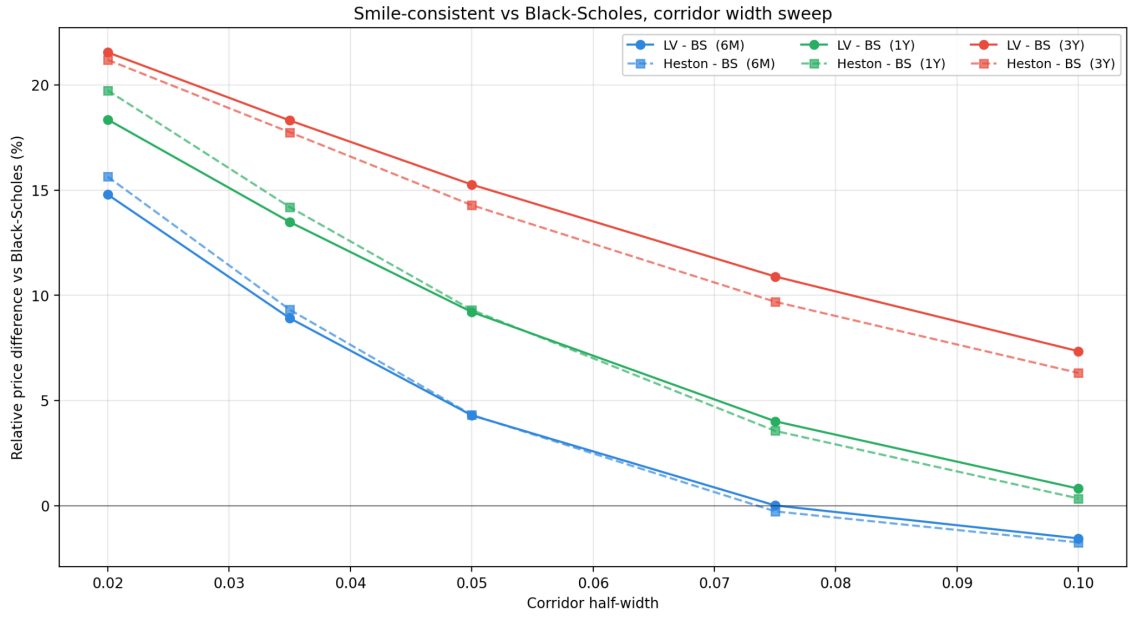

Figure 4.3: Relative price differences of the local-volatility and multi-factor Heston models with respect to the Black–Scholes benchmark, expressed as a percentage of the benchmark price, as a function of corridor half-width. Solid lines show (LV − BS)/BS and dashed lines show (Heston − BS)/BS. Colours indicate maturity.

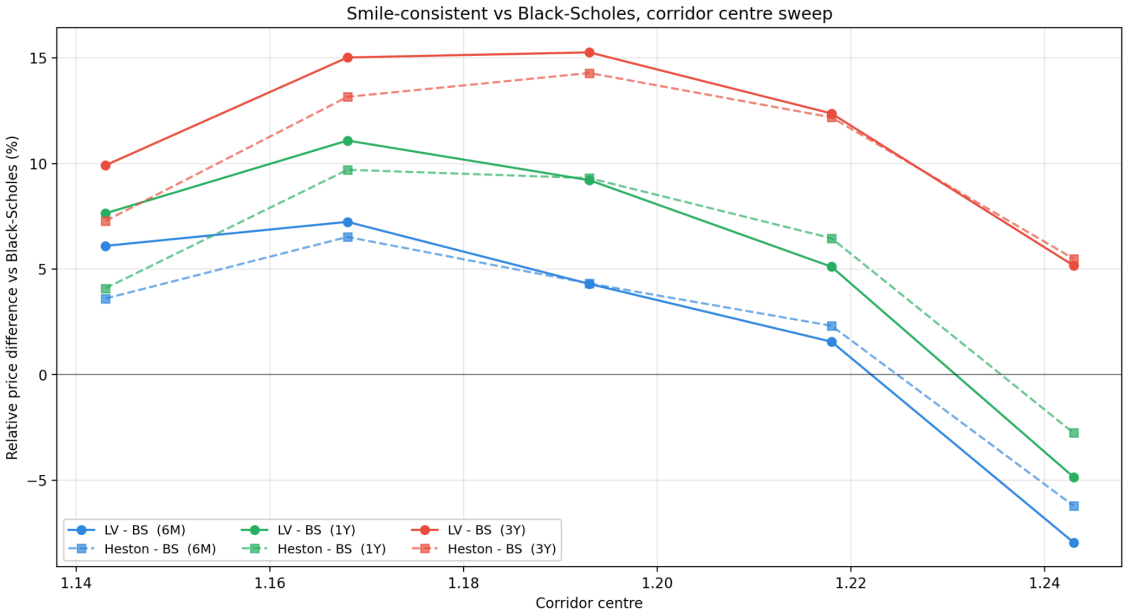

Figure 4.4: Relative price differences of the local-volatility and multi-factor Heston models with respect to the Black–Scholes benchmark, expressed as a percentage of the benchmark price, as a function of corridor centre. The half-width is fixed at 0.05, and only the central part of the centre sweep is shown. Solid lines show (LV−BS)/BS and dashed lines show (Heston − BS)/BS. Colours indicate maturity.

For the width sweep, the gap between the smile-consistent models and the benchmark is largest for narrow corridors and decreases as the corridor becomes wider. At a half-width of 0.02, the local-volatility and Heston prices are about 14.6–14.8% above Black–Scholes at 6M, about 18.3–19.1% above at 1Y, and about 21.3% above at 3Y. At a half-width of 0.10, the difference is much smaller: about −1.6 to −1.7% at 6M, about 0.4–0.7% at 1Y, and about 6.7–7.2% at 3Y.

This shows that the volatility smile matters most for narrow corridors. This can be understood intuitively from the payoff structure. A narrow corridor depends heavily on the probability mass located close to the corridor boundaries. The volatility smile alters the shape of the risk-neutral distribution precisely in these regions, through skew and fat-tail effects. As a result, small changes in the local density translate into large relative changes in the corridor probability. In contrast, a wider corridor captures a larger fraction of the distribution, making the price less sensitive to local distortions introduced by the smile. In the limiting case of an infinitely wide corridor, all models would converge to the discounted coupon value.

For the centre sweep, the gap is asymmetric in the corridor centre. The smile-consistent models are generally above Black–Scholes when the corridor is near the spot or slightly below it, although the two models are not identical at every centre. At a centre offset of −0.025, the local-volatility and Heston prices are about 6.4–7.2% above Black–Scholes at 6M, about 9.7–10.9% above at 1Y, and about 13.6–14.7% above at 3Y. At the spot-centred corridor, the differences are about 4.1–4.4%, about

9.2%, and about 14.7–15.2% for 6M, 1Y, and 3Y respectively.

When the corridor is shifted upward, the difference becomes smaller and can change sign at the shorter maturities. At a centre offset of +0.05, the local-volatility and Heston prices are about 6.4–7.8% below Black–Scholes at 6M and about 2.9–4.8% below at 1Y. At 3Y, both smile-consistent prices remain above Black–Scholes by about 5.2–5.8%.

The same figures also show the relationship between the local-volatility and multi-factor Heston prices. In the width sweep, the solid and dashed lines follow almost the same decreasing pattern as the corridor becomes wider, and the two smile-consistent models are close in absolute price. In the centre sweep, the broad asymmetric shape is also similar, but pointwise differences between local volatility and Heston are more visible, especially away from the spot-centred corridor.

The LV-Heston comparison should therefore not be interpreted as a strict pointwise ranking between the two smile-consistent models. Both models are calibrated to the same set of implied-volatility smiles and produce the same main pricing message relative to the Black-Scholes benchmark, but their dynamics are different and this can matter for individual corridor locations. The range-accrual payoff depends on the distribution of the underlying at several fixing dates, so differences between the local-volatility and stochastic-volatility dynamics can still appear in the centre sweep even when the overall price patterns are similar.

A caveat applies to the direct comparison between the two smile-consistent models. The local-volatility model reproduces the fitted implied-volatility surfaces by construction through the Dupire formula, while the multi-factor Heston model fits these surfaces with an overall RMSE of about 10 basis points. Part of the pointwise difference between the local-volatility and Heston prices therefore reflects calibration residual on the Heston side rather than a structural difference between deterministic and stochastic volatility dynamics. This is most relevant for narrow corridors and for corridor centres far from the spot, where the price depends sensitively on probability mass located in regions where the calibration fit is least tight.

The main conclusion from these results is that the smile effect is corridor dependent. For some corridors, the smile-consistent models give higher prices than the at-the-money Black–Scholes benchmark, while for others they give lower prices. The difference depends on both the corridor width and the corridor centre. Thus, within the quanto-adjusted pricing framework, the at-the-money volatility term structure is not sufficient to describe the corridor probabilities accurately. The two smile-consistent models give similar qualitative conclusions, although their pointwise prices are not identical across all corridor locations.

5

Conclusion and outlook

This thesis investigated how volatility modelling affects the valuation of a EUR/USD range accrual quantoed to GBP. Three pricing models were compared: an at-the-money Black–Scholes benchmark, a triangle-based Dupire local-volatility model, and a multi-factor Heston model of De Col, Gnoatto, and Grasselli. The models were applied to the same market data, discount curves, fixing schedule, and GBP pricing measure, so that the comparison isolates the effect of the volatility modelling assumptions.

The main conclusion is that the volatility smile has a material effect on the range-accrual price within the quanto-adjusted pricing framework. The Black–Scholes benchmark captures the quanto adjustment and the at-the-money volatility term structure, but it does not use strike-dependent implied volatilities. The differences between the benchmark and the local-volatility and multi-factor Heston models therefore reflect the effect of smile-consistent volatility modelling. The effect was strongest for narrow corridors, where the price is highly sensitive to the probability mass close to the corridor boundaries. For wider corridors, the difference became smaller, since a wider corridor captures a larger part of the distribution. In the limit of a sufficiently wide corridor, all models should converge to the discounted coupon value.

The centre sweep showed that the smile effect is not a simple constant correction to the Black–Scholes price. The smile-consistent models produce a higher central peak than the at-the-money benchmark, especially for longer maturities, reflecting a redistribution of probability mass under the implied-volatility smile. However, the ordering between the models changes across the sweep. In parts of the intermediate off-centre region, Black–Scholes produces higher prices than the smile-consistent models, while the smile-consistent prices increase relative to the benchmark again in parts of the tails. This reflects how the smile reshapes the risk-neutral distribution, shifting mass between the centre, intermediate region, and tails. The effect of the volatility smile is therefore inherently corridor dependent: both the magnitude and the sign of the pricing correction depend on the corridor width and its location relative to the spot.

The local-volatility and multi-factor Heston models give similar qualitative conclusions, but their pointwise prices are not identical across all corridor locations. In the width sweep, the two smile-consistent models are close in absolute price and follow almost the same pattern as the corridor becomes wider. In the centre sweep,

some larger pointwise differences appear away from the spot-centred corridor. The comparison should therefore not be interpreted as a strict ranking between local volatility and Heston. Both models are calibrated to reproduce the same implied-volatility smiles and produce the same main conclusion relative to the Black-Scholes benchmark: smile-consistent volatility modelling changes the corridor probabilities in a material and corridor-dependent way. Part of the residual pointwise difference between the two smile-consistent models also reflects calibration residual on the Heston side, since the local-volatility construction reproduces the fitted smiles exactly while the multi-factor Heston model fits them to about 10 basis points overall.

Natural extensions would be to repeat the analysis across several market dates and to test different corridor grids. It would also be useful to apply the framework to other currency triangles, especially those involving less commonly traded currencies. Less liquid pairs would test the robustness of the calibration when smile data are sparse or noisy. Another extension would be to compare hedge sensitivities, since two models can produce similar prices but different risk exposures. A further natural step would be to consider a stochastic local-volatility specification, which combines the exact smile fit of the Dupire construction with the stochastic variance dynamics of the multi-factor Heston model through a leverage function applied to the stochastic-volatility backbone. Such a model would remove the calibration residual that affects the multi-factor Heston comparison while retaining stochastic volatility.

Overall, the results show that the quanto adjustment and the volatility smile enter the pricing problem in different ways. The quanto adjustment determines the appropriate GBP-measure dynamics for the EUR/USD spot, while the volatility smile affects the distribution of EUR/USD at the fixing dates. The numerical comparison in this thesis is therefore conditional on the quanto-adjusted pricing framework. Within this framework, replacing an at-the-money volatility benchmark by smile-consistent volatility models has a material and corridor-dependent effect on the range-accrual price. The local-volatility and multi-factor Heston specifications give similar qualitative conclusions, but their pointwise differences can still matter for individual corridor locations and for risk-management quantities.

Bibliography

[1] Mark B. Garman and Steven W. Kohlhagen. "Foreign Currency Option Values". In: *Journal of International Money and Finance* 2.3 (1983). DOI: 10.1016/S0261-5606(83)80001-1.

[2] Iain J. Clark. *Foreign Exchange Option Pricing: A Practitioner's Guide*. Wiley, 2011. DOI: 10.1002/9781119208679.

[3] Bruno Dupire. "Pricing with a Smile". In: *Risk* 7.1 (1994).

[4] Jim Gatheral. *The Volatility Surface. A Practitioner's Guide*. Wiley, 2006.

[5] Alvise De Col, Alessandro Gnoatto, and Martino Grasselli. *Smiles all around: FX joint calibration in a multi-Heston model*. arXiv:1201.1782v3. 2013. arXiv: 1201.1782 [q-fin.PR].

[6] Fima C. Klebaner. *Introduction to Stochastic Calculus with Applications*. 3rd ed. London: Imperial College Press, 2012. DOI: 10.1142/P821.

[7] Steven E. Shreve. *Stochastic Calculus for Finance II: Continuous-Time Models*. New York: Springer, 2004. ISBN: 9780387401010.

[8] Uwe Wystup. *Foreign Exchange Quanto Options*. Working Paper 10. Frankfurt School of Finance & Management, Centre for Practical Quantitative Finance, 2008.

[9] Matthias R. Fengler. "Option Data and Modeling BSM Implied Volatility". In: *Handbook of Computational Finance*. Ed. by Jin-Chuan Duan, James E. Gentle, and Wolfgang Karl Härdle. Berlin and Heidelberg: Springer, 2012. Chap. 6, pp. 117–153. DOI: 10.1007/978-3-642-17254-0_6.

[10] Julien Guyon. *A New Class of Local Correlation Models*. SSRN Electronic Journal. SSRN 2283419. 2013.

[11] Rama Cont. "Volatility Clustering in Financial Markets: Empirical Facts and Agent-Based Models". In: *Long Memory in Economics*. Ed. by Gilles Teyssière and Alan P. Kirman. Berlin, Heidelberg: Springer, 2007, pp. 289–309. DOI: 10.1007/978-3-540-34625-8_10.

[12] Paul Glasserman. *Monte Carlo Methods in Financial Engineering*. Vol. 53. Applications of Mathematics. New York: Springer, 2003. DOI: 10.1007/978-0-387-21617-1.

[13] Roger Lord, Remmert Koekkoek, and Dick van Dijk. "A Comparison of Biased Simulation Schemes for Stochastic Volatility Models". In: *Quantitative Finance* 10.2 (2010). DOI: 10.1080/14697680802392496.

A

Calibration diagnostics

This appendix provides the calibration diagnostics underlying the summary in Section 4.1. The figures document the SVI smile fits, the fitted implied-volatility surfaces, the extracted local-volatility surfaces, and the multi-factor Heston calibration diagnostics. They are included here to show the numerical quality of the calibration and surface construction without interrupting the main price-comparison discussion in Chapter 4.

A.1 Implied-volatility and local-volatility surfaces

Figures A.1–A.3 show the fitted implied-volatility surfaces and the corresponding local-volatility surfaces for EUR/USD, GBP/USD, and EUR/GBP. The implied-volatility surfaces are obtained from the SVI total-variance construction described in Section 3.4. The local-volatility surfaces are then extracted using Dupire's formula in the Gatheral form, with the numerical safeguards described in Section 3.4.

The implied-volatility surfaces are generally smoother than the local-volatility surfaces because the Dupire extraction depends on derivatives of the fitted total-variance surface. The local-volatility surfaces therefore provide a useful diagnostic of the numerical stability of the surface construction.

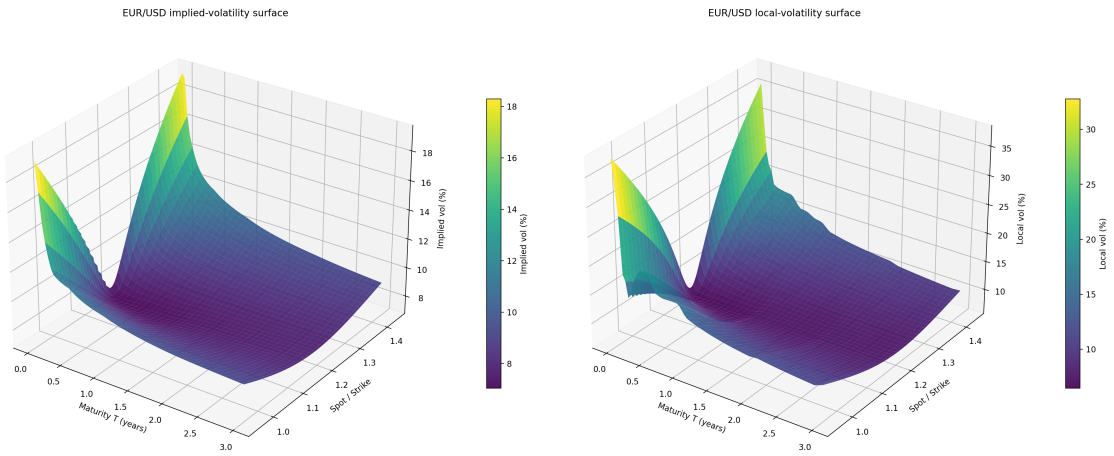

(a) Implied-volatility surface. **(b)** Local-volatility surface.

Figure A.1: EUR/USD fitted implied-volatility surface and extracted local-volatility surface.

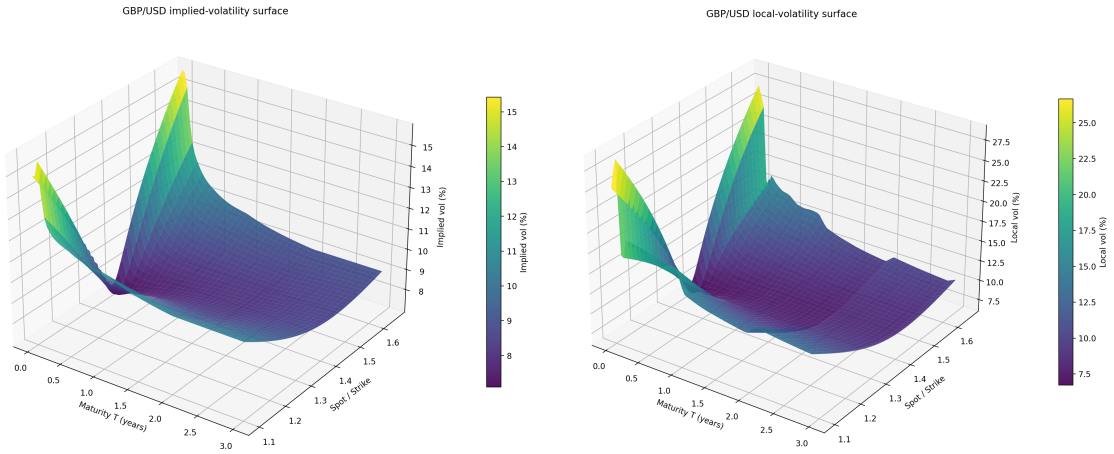

(a) Implied-volatility surface. **(b)** Local-volatility surface.

Figure A.2: GBP/USD fitted implied-volatility surface and extracted local-volatility surface.

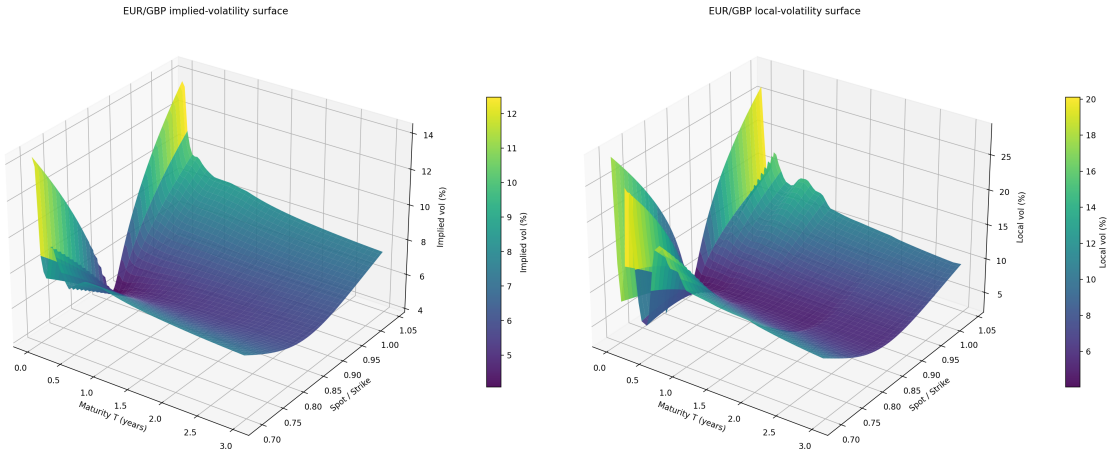

(a) Implied-volatility surface. **(b)** Local-volatility surface.

Figure A.3: EUR/GBP fitted implied-volatility surface and extracted local-volatility surface.

A.2 Multi-factor Heston calibration diagnostics

This subsection reports the detailed calibration diagnostics for the multi-factor Heston model. The calibration is performed jointly across EUR/USD, GBP/USD, and EUR/GBP at the maturities 6M, 1Y, and 3Y, using the procedure described in Section 3.5. The main text reports the resulting RMSE values in Table 4.1. The figures below provide the corresponding visual diagnostics.

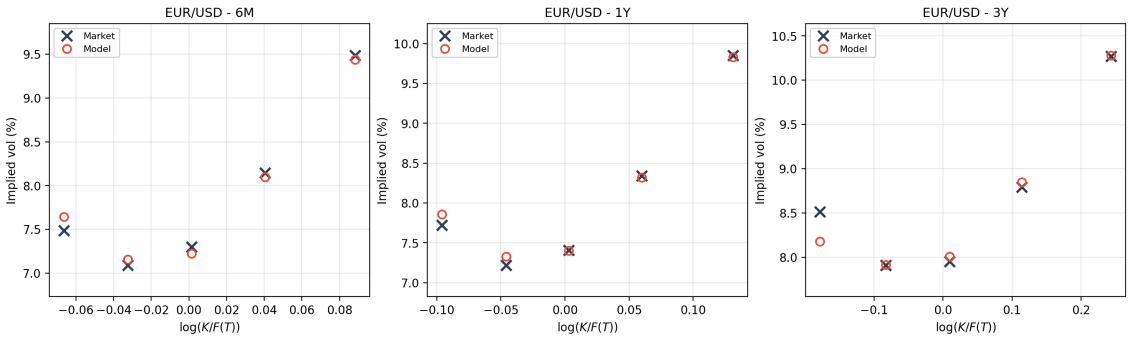

Figure A.4: Multi-factor Heston fit to the EUR/USD implied-volatility smile across the calibrated maturities. Market implied volatilities are compared with model-implied volatilities obtained by Fourier pricing and Garman–Kohlhagen implied-volatility inversion.

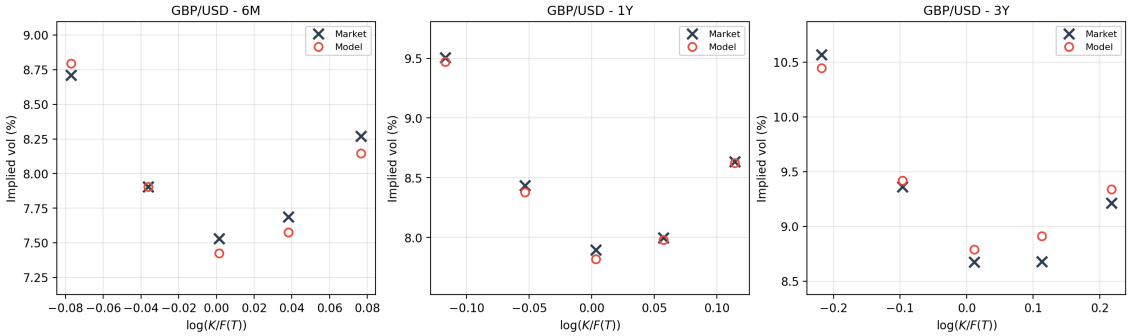

Figure A.5: Multi-factor Heston fit to the GBP/USD implied-volatility smile across the calibrated maturities. Market implied volatilities are compared with model-implied volatilities obtained by Fourier pricing and Garman–Kohlhagen implied-volatility inversion.

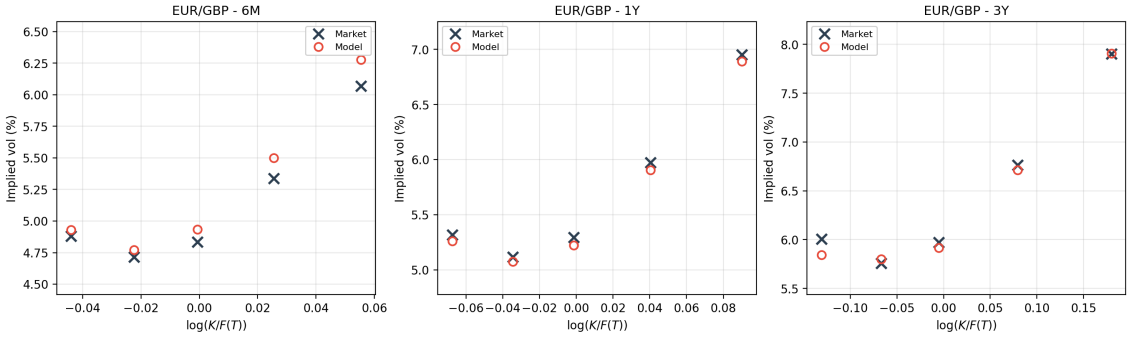

Figure A.6: Multi-factor Heston fit to the EUR/GBP implied-volatility smile across the calibrated maturities. Market implied volatilities are compared with model-implied volatilities obtained by Fourier pricing and Garman–Kohlhagen implied-volatility inversion.

The fit is close across the three currency pairs, with the main residual errors occurring in some wing quotes. This is consistent with the calibration summary in Section

4.1, where the overall implied-volatility RMSE was reported as 10.39 basis points. The appendix diagnostics therefore support the use of the calibrated multi-factor Heston model in the price comparison.

DEPARTMENT OF MATHEMATICAL SCIENCES
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se

CHALMERS
UNIVERSITY OF TECHNOLOGY