# OIS Cap Pricing using a CIR++ Model with Time-Dependent Volatility

Viktor Pitkänen Fungbrant

Master's thesis
2026:E24

Lund Institute of Technology
Lund University

Faculty of Engineering
Centre for Mathematical Sciences
Mathematical Statistics

Centrum Scientiarum Mathematicarum

Master's Theses in Mathematical Sciences 2026:E24
ISSN 1404-6342

LUTFMS-3552-2026

Mathematical Statistics
Centre for Mathematical Sciences
Lund University
Box 118, SE-221 00 Lund, Sweden

http://www.maths.lu.se/

# Abstract

This thesis studies the pricing of overnight indexed swap (OIS) caps using a one-factor CIR++ short-rate model with piecewise constant volatility. The model combines a CIR process with a deterministic shift, allowing an exact fit to the initial zero-coupon term structure and thereby enabling the remaining CIR parameters to be calibrated to other market instruments. OIS caplets are priced using Laplace transform based methods, and cap prices are obtained as sums of the underlying caplets.

The empirical analysis evaluates the model's performance in two distinct market environments characterized by different yield-curve shapes. The study examines in-sample fit across maturities, out-of-sample pricing performance, pricing across strikes, and the effect of using different calibration strike sets and calibration windows. The results show that the model reproduces observed market cap prices well in sample and captures the cap price surface across strikes and maturities with reasonable accuracy. Out-of-sample performance is weaker than in sample, but remains satisfactory overall. The estimated volatility parameters appear relatively stable across maturities and over time. Overall, the findings suggest that the CIR++ model with piecewise constant volatility provides a tractable and reasonably accurate approach for price-based calibration and pricing of OIS caps, although some limitations remain, particularly regarding implied-volatility inversion for certain contracts.

## Keywords

OIS caps, SOFR, CIR++ model, deterministic shift, time-dependent volatility, affine term-structure models, Laplace transform methods, overnight rates

## Acknowledgment

I would like to thank my supervisor, Magnus Wiktorsson, for valuable guidance and always taking the time to help me throughout the thesis process. Your support and feedback have been greatly appreciated.

# Contents

# 1 Background

Interest rate derivatives (IRDs) are fundamental to the financial markets. They are widely used to hedge interest rate risk, support risk management, and take positions on future interest rate movements. IRDs also constitute the largest segment of the global over-the-counter (OTC) derivatives market by notional amount outstanding. According to the Bank for International Settlements (BIS), the notional value of outstanding OTC derivatives were USD 846 trillion at June 2025, of which interest rate derivatives accounted for 79% (Bank for International Settlements (2025a)). In terms of trading activity, OTC interest rate derivatives had an average daily turnover of USD 7.9 trillion in April 2025, with overnight index swaps (OIS) being the most traded instrument, averaging at USD 5.1 trillion per day (Bank for International Settlements (2025b)). These figures illustrate both the scale of the market and the importance of understanding how such instruments are priced.

For many years the London Interbank Offered Rate (LIBOR) played a major role in the global financial system. It acted as a reference rate for financial contracts while also functioning as a benchmark against which financing costs were measured. LIBOR represented the rate at which the world's largest and most financially sound institutions were able to borrow from one another on a short-term basis. It was determined daily based on a panel of banks estimates of the rate at which they could barrow from other banks in a reasonable market size. It was later revealed that several banks had manipulated their reported estimates, which contributed to the decision to phase out LIBOR (Huan et al. (2023)).

LIBOR has gradually been replaced by alternative nearly risk-free rates (RFRs), which are based on actual transactions rather than subjective estimates. In the United States, the relevant reference rate is the Secured Overnight Financing Rate (SOFR), which is a broad measure of the cost of borrowing cash overnight collateralized by U.S. Treasury securities in the repurchase agreement market (Federal Reserve Bank of New York (nd)). The transition from LIBOR to RFRs has led to substantial changes in interest rate derivative markets. A key difference is that LIBOR was forward-looking, meaning that the rate for an accrual period was known at its beginning, whereas SOFR and other RFRs are overnight rates that are backward-looking and known at the end of the accrual period. This has important implications for both discounting and the pricing of interest rate derivatives (Lyashenko and Mercurio (2019)). An overnight index swap is an interest rate swap in which the floating leg is determined by a compounded overnight rate, making the OIS market especially relevant in the post-LIBOR environment.

Against this background, the purpose of this thesis is to study the pricing of OIS caps using a CIR++ short-rate model with piecewise constant volatility. The study builds on earlier work by Stoltz and Sönne (2025), as it addresses the same problem of pricing OIS caps, but uses a different modeling approach. The model is calibrated to market data and evaluated with respect to its in-sample fit, its ability to reproduce prices across strikes and maturities, and its out-of-sample performance under different market conditions. In this way, the thesis examines how well a relatively tractable short-rate model can capture observed prices in the SOFR OIS cap market.

# 2 Theory

This section presents the theoretical framework underlying the study. It introduces the basic concepts and mathematics required for modeling the short rate and pricing interest rate derivatives. The presentation begins with general interest rate theory and then moves on to the more specialized theory relevant for the modeling and pricing approach considered in this thesis.

## 2.1 Bank Account

The bank account is fundamental in interest rate theory and serves as a risk-free investment where profit is accrued continuously at the stochastic risk-free rate. The definition of a bank account is stated below (Brigo and Mercurio (2006)).

**Definition 1** (Bank account)**.** *The value of the bank account is defined as $B(t)$ for $t \geq 0$ and it is assumed that $B(0) = 1$. The dynamics of the bank account are given by the differential equation stated below*

$$dB(t) = r(t)B(t)dt, \; B(0) = 1 \tag{1}$$

*where $r(t)$ is a positive function of time. The solution to this equation subject to the initial condition is*

$$B(t) = e^{\int_0^t r(s)ds}. \tag{2}$$

### 2.1.1 Discounting

The definitions above show that a unit amount deposited at time $t = 0$ will be worth $B(t)$ at time t. Consequently, an amount $M_0$ deposited at time 0 is worth $M_0B(t)$ at time t. Next consider how much that has to be deposited at time 0 to make the value 1 at time T

$$M_0B(T) = 1. \tag{3}$$

Assuming for now that the interest rate $r$ is deterministic yields that

$$M_0 = \frac{1}{B(T)}. \tag{4}$$

The value of depositing this amount of currency at time 0 yields a value at time t which is

$$M_0B(t) = \frac{B(t)}{B(T)}. \tag{5}$$

This is thus the corresponding value at time t of one unit of currency at time T.

Now, if the interest rate $r$ is assumed to be stochastic again the stochastic discount factor can be defined (Brigo and Mercurio (2006)).

**Definition 2** (Discount factor)**.** *The Discount factor $D(t,T)$ between $t$ and $T$ is the amount of currency at time $t$ that results in one unit of currency at time $T$*

$$D(t,T) = \frac{B(t)}{B(T)} = e^{-\int_t^T r(s)ds}. \tag{6}$$

## 2.2 No-Arbitrage Pricing

A fundamental concept in economics is the absence of arbitrage. An arbitrage opportunity is defined as the possibility of making a strictly positive profit without investing anything or taking any risk (Björk (2009)). In well-functioning financial markets, such opportunities should not exist.

The ability to price contingent claims, i.e. financial contracts whose payoff is determined at maturity $T$, in a no-arbitrage framework is therefore of central importance. Mathematically, the absence of arbitrage is characterized by the existence of an equivalent martingale measure (also referred to as a risk-neutral measure). Assuming the existence of an equivalent martingale measure $\mathbb{Q}$ and letting $H_T$ denote the payoff of an attainable contingent claim at maturity $T$, which is a contingent claim that can be replicated by a self-financing trading strategy, the unique arbitrage-free price $\Pi(t)$ associated with $H$ at time $t$ is given by (Brigo and Mercurio (2006))

$$\Pi(t) = E^{\mathbb{Q}}\left[D(t,T)H|\mathcal{F}_t\right] \tag{7}$$

where $D(t,T)$ is the discount factor. $\{\mathcal{F}_t\}_{t\geq 0}$ is denoted the filtration at time $t$ and models the flow of information. $\mathcal{F}_t$ represents all the information available at time $t$.

## 2.3 Zero-coupon Bonds

**Definition 3** (Zero-coupon Bond)**.** *The zero-coupon bond is a contract which guarantees the holder one unit of currency at time $T$. The price of this contract at $t < T$ is denoted $P(t,T)$ (Brigo and Mercurio (2006)).*

The price $P(t,T)$ represents the value at time $t$ of the future payment of a unit of currency at $T$. A zero-coupon bond can be viewed as a contingent claim with the deterministic payoff $H_T = 1$ at maturity. Therefore, using expression (7), the price of the zero-coupon bond under the risk-neutral measure $\mathbb{Q}$ is

$$P(t,T) = E^{\mathbb{Q}}\left[e^{-\int_t^T r(s)ds}|\mathcal{F}_t\right] \tag{8}$$

where $\mathcal{F}_t$ is the filtration at time $t$.

## 2.4 Yield

**Definition 4** (Yield)**.** *The yield, denoted $R(t,T)$ is defined as the constant rate at which an investment of $P(t,T)$ units of currency at time $t$ accrues continuously to yield a unit amount of currency at maturity $T$. The formula for the yield is given by*

$$R(t,T) := -\frac{\ln P(t,T)}{\tau} \tag{9}$$

*where $\tau$ is the time to maturity (Brigo and Mercurio (2006)).*

Time to maturity is defined as $\tau = T - t$, representing the time difference between time $t$ and maturity $T$. See section 3.1 for more information about how this is implemented in practice, particularly in this study. With this expression, the zero-coupon-bond price can also be written as

$$P(t,T) = e^{-R(t,T)\tau}. \tag{10}$$

The yield can therefore be seen as the constant continuously compounded rate that reproduces the price of a zero-coupon bond.

In this report, references to the zero-coupon curve or the yield curve denote the mapping

$$T \mapsto R(t,T),\ T > t \tag{11}$$

The shape of the yield curve can provide information about the market's expectation of future interest rates and the state of the economy. Under normal market conditions the yield curve typically exhibits an upward sloping shape where yields are higher for longer maturities. This can be interpreted as the market expecting higher future rates and economic growth. An inverted yield curve is represented by short-term bonds having higher yields than long-term bonds. This can be interpreted to mean that future interest rates are expected to decrease and that a recession is near or already ongoing.

## 2.5 Forward Rates

The forward rate will be derived as the fixed interest rate that turns the price of a forward rate agreement (FRA) into a fair contract. As a first step the FRA is defined.

**Definition 5** (Forward rate agreement (FRA))**.** *The FRA involves three time instants, the current time $t$, expiry $S$ and maturity $T$. The contract works by exchanging a fixed payment based on a fixed rate $K$, known at $t$, against a floating payment based on the spot rate resetting in $S$ and with maturity $T$. The interest rates are simply compounded. The value of the contract at time $t$ is*

$$V_{FRA} = N[P(t,T)\tau K - P(t,S) + P(t,T)] \tag{12}$$

*where $N$ is the nominal value and $\tau$ is the year fraction between $S$ and $T$ (Brigo and Mercurio (2006)).*

To get the fair value of this contract, formula (12) is set equal to zero and solved for $K$, which defines the simply compounded forward rate.

$$\begin{gathered} N[P(t,T)\tau K - P(t,S) + P(t,T)] = 0 \\ P(t,T)\tau K = P(t,S) - P(t,T) \\ K = \frac{P(t,S) - P(t,T)}{P(t,T)\tau} = \frac{1}{\tau}\left(\frac{P(t,S)}{P(t,T)} - 1\right) \coloneqq F(t;S,T) \end{gathered} \tag{13}$$

By letting the maturity approach expiry, $T \to S^+$, the instantaneous forward rate can be defined (Brigo and Mercurio (2006)).

**Definition 6** (Instantaneous Forward Rate)**.** *The instantaneous forward rate $f(t,T)$ at time $t$ with maturity $T > t$ is defined as*

$$f(t,T) \coloneqq \lim_{T \to S^+} F(t;S,T) = -\frac{\partial \ln P(t,S)}{\partial S}. \tag{14}$$

Using the instantaneous forward rate, the price of the zero-coupon bond can also be written as

$$P(t,T) = e^{-\int_t^T f(t,u)du}. \tag{15}$$

## 2.6 One-Factor Short-Rate Models with Deterministic Shifts

In this study a stochastic one-factor model with deterministic shift is used to model the short rate, as shown in (16). This type of model combines a single-factor stochastic process that drives the dynamics of the short rate with a deterministic function that makes the model fit exactly to the initial term structure of interest rates. In expression (16) $r_t$ is the short rate, $X_t$ is the stochastic one-factor process and $\varphi_t$ is the deterministic function.

$$r_t = X_t + \varphi_t \tag{16}$$

The function $\varphi$ is defined such that

$$p_{\varphi}(t_c, T) = e^{-\int_{t_c}^{T} \varphi(s;t_c)\ ds} = \frac{p_{mk}(t_c, T)}{p_X(t_c, T)} \tag{17}$$

where $P_{mk}(t_c, T)$ is the observed market zero-coupon curve and $t_c$ stands for the point in time of calibration to market data. By this definition of $\varphi$ the short rate model is able to exactly fit to the initial term structure of interest rates, see below.

$$\begin{aligned} p_r(t_c, T) &= E^{\mathbb{Q}}\left[e^{-\int_{t_c}^{T} r(s)\ ds} | \mathcal{F}_{t_c}\right] \\ &= E^{\mathbb{Q}}\left[e^{-\int_{t_c}^{T} X(s)\ ds} | \mathcal{F}_{t_c}\right] e^{-\int_{t_c}^{T} \varphi(s;t_c)\ ds} \\ &= p_X(t_c, T)\frac{p_{mk}(t_c, T)}{p_X(t_c, T)} \\ &= p_{mk}(t_c, T) \end{aligned} \tag{18}$$

## 2.7 Affine Term-Structure Models

Interest rate models that can be written in the form (19) are called affine term-structure models.

$$dr(t) = (\alpha(t)r(t) + \beta(t))dt + \sqrt{\gamma(t)r(t) + \delta(t)}dW(t) \tag{19}$$

$r(t)$ is the short-rate process and $W(t)$ is a Brownian motion (Alvarez (1998)). Models of this form all possess the property thar the price of a zero-coupon bond can be written in the form

$$P(t, T) = e^{A(t,T) - B(t,T)r(t)}. \tag{20}$$

$A(t, T)$ and $B(t, T)$ are two functions that satisfy the following system of ordinary differential equations

$$\begin{aligned} B_t'(t, u) &= -\alpha(t)B(t, u) + \frac{1}{2}\gamma(t)B^2(t, u) - 1,\ B(u, u) = 0, \\ A_t'(t, u) &= \beta(t)B(t, u) - \frac{1}{2}\delta(t)B^2(t, u),\ A(u, u) = 0. \end{aligned} \tag{21}$$

where $\alpha(t), \gamma(t), \beta(t)$ and $\delta(t)$ come from (19) (Alvarez (1998)).

## 2.8 Overnight Rate

Overnight rates are interest rates that apply to borrowing from one day to the next. Since these loans have a very short term, there's little credit risk associated with them. In addition, many of the loans that these rates are based on are also secured, further decreasing the risk. For these reasons overnight rates are classified as RFRs. Overnight rates can be compounded over time to construct rates with longer terms. The compounded overnight rate over the accrual period $\tau_i$ between $T_{i-1}$ and $T_i$ is given by

$$\prod_{k=0}^{n-1} \frac{1}{p_r(t_k, t_k + \Delta)} = e^{\sum_{k=0}^{n-1} \int_0^{\Delta} f(t_k, t_k+s)ds} \tag{22}$$

where $T_{i-1} = t_0 < t_1 < ... < t_{n-1} + \Delta = t_n = T_i$ and $\Delta$ is the year fraction corresponding to one day. A common way to approximate (22) is as below.

$$\begin{aligned} e^{\sum_{k=0}^{n-1} \int_0^{\Delta} f(t_k, t_k+s)ds} &\approx e^{\sum_{k=0}^{n-1} f(t_k,t_k)\Delta} \\ &= e^{\sum_{k=0}^{n-1} r(t_k)\Delta} \\ &\approx e^{\sum_{k=0}^{n-1} \int_{t_k}^{t_k+\Delta} r(s)ds} \\ &= e^{\int_{T_{i-1}}^{T_i} r(s)ds} \end{aligned} \tag{23}$$

This is the continuously compounded short rate.

## 2.9 Overnight Index Swap

An interest rate swap is a contract where a payment stream at a fixed interest rate, known as the swap rate, is exchanged for a payment stream at a floating rate. The OIS is an interest rate swap where the floating rate is the compounded overnight rate over the accrual period.

The swap rate is determined such that the value of the contract is zero when it is made. Let $\{T_0, T_1, ..., T_n\}$ be a set of dates where $T_0 = 0$, the time when the contract is written, and $\{T_1, ..., T_n\}$ are the payment dates. The spacing between the dates are defined as $\tau_i = T_i - T_{i-1}$. The expression for the swap rate is then stated as below, see (Björk (2009)) for more information.

$$R = \frac{1 - P(T_0, T_n)}{\sum_{i=1}^{n} \tau_i P(T_0, T_i)} \tag{24}$$

$P(T_0, T_i)$ is the price at time $T_0$ of a zero-coupon bond maturing at $T_i$.

## 2.10 OIS Caplet

The OIS caplet is a contract that mainly serves as a protection against rising interest rates limiting the cost for those that have to pay floating rates. An OIS caplet is a European call option on the compounded overnight rate. For a period of time $\tau_i = T_i - T_{i-1}$ the payoff is

$$\left( \prod_{k=0}^{n-1} \frac{1}{p_r(t_k, t_k + \Delta)} - (1 + K\tau_i) \right)^+ \tag{25}$$

where $T_{i-1} = t_0 < t_1 < ... < t_n = T_i$ and $\Delta = t_{k+1} - t_k$ is the year fraction equivalent to one day. Using the approximation for the compounded overnight rate (23) the payoff can be expressed as below.

$$\left( e^{\int_{T_{i-1}}^{T_i} r(s)ds} - (1 + K\tau_i) \right)^+ \tag{26}$$

To price this contract in a way consistent with no arbitrage, the risk-neutral expectation needs to be calculated.

$$\begin{aligned} &E^{\mathbb{Q}} \left[ e^{-\int_t^{T_i} r(s)ds} \left( e^{\int_{T_{i-1}}^{T_i} r(s)ds} - (1 + K\tau_i) \right)^+ |\mathcal{F}_t \right] = \\ &= E^{\mathbb{Q}} \left[ (1 + K\tau_i) e^{-\int_t^{T_{i-1}} r(s)ds} \left( \frac{1}{1 + K\tau_i} - e^{-\int_{T_{i-1}}^{T_i} r(s)ds} \right)^+ |\mathcal{F}_t \right] \end{aligned} \tag{27}$$

Using the notation in (17) equation (27) can be written as

$$
\begin{aligned}
& E^{\mathbb{Q}}\left[(1+K\tau_i)e^{-\int_t^{T_{i-1}} X(s)ds}e^{-\int_t^{T_{i-1}} \varphi(s)ds}\left(\frac{1}{1+K\tau_i}-e^{-\int_{T_{i-1}}^{T_i} X(s)ds}e^{-\int_{T_{i-1}}^{T_i} \varphi(s)ds}\right)^+ |\mathcal{F}_t\right] = \\
& = E^{\mathbb{Q}}\left[(1+K\tau_i)p_\varphi(t,T_{i-1})e^{-\int_t^{T_{i-1}} X(s)ds}\left(\frac{1}{1+K\tau_i}-e^{-\int_{T_{i-1}}^{T_i} X(s)ds}p_\varphi(T_{i-1},T_i)\right)^+ |\mathcal{F}_t\right] = \\
& = E^{\mathbb{Q}}\left[(1+K\tau_i)p_\varphi(t,T_{i-1})e^{-\int_t^{T_{i-1}} X(s)ds}\left(\frac{1}{1+K\tau_i}-e^{-\int_{T_{i-1}}^{T_i} X(s)ds}\frac{p_\varphi(t,T_i)}{p_\varphi(t,T_{i-1})}\right)^+ |\mathcal{F}_t\right] = \\
& = E^{\mathbb{Q}}\left[(1+K\tau_i)p_\varphi(t,T_i)e^{-\int_t^{T_{i-1}} X(s)ds}\left(\frac{p_\varphi(t,T_{i-1})}{(1+K\tau_i)p_\varphi(t,T_i)}-e^{-\int_{T_{i-1}}^{T_i} X(s)ds}\right)^+ |\mathcal{F}_t\right].
\end{aligned}
\tag{28}
$$

By applying the tower property of conditional expectations, expression (28) can also be written as

$$
E^{\mathbb{Q}}\left[(1+K\tau_i)p_\varphi(t,T_i)e^{-\int_t^{T_{i-1}} X(s)ds}E^{\mathbb{Q}}\left[\left(\frac{p_\varphi(t,T_{i-1})}{(1+K\tau_i)p_\varphi(t,T_i)}-e^{-\int_{T_{i-1}}^{T_i} X(s)ds}\right)^+ |\mathcal{F}_{T_{i-1}}\right] |\mathcal{F}_t\right]. \tag{29}
$$

This representation separates the caplet price into a part that is known at the beginning of the accrual period and a remaining stochastic part. To derive the pricing formula used in this study, an approach based on the Laplace transform is applied to expression (29). Further details are provided in section 2.12.

## 2.11 Cap

The cap is a financial contract that consists of a series of caplets. For a cap consisting of $n$ caplets with tenor structure $\bar{T} = [T_0, T_1, T_2, ..., T_{n-1}, T_n]$ the price is equal to the sum of the caplet prices.

$$
P_{cap}(t,\bar{T}) = \sum_{k=1}^{n} P_{caplet}(t, T_{k-1}, T_k) \tag{30}
$$

### 2.11.1 Calculation of At-The-Money Strikes for Caps

In order to identify at-the-money (ATM) caps, the strike that makes the value of the corresponding sum of FRAs zero at initiation must be calculated. For a cap consisting of n caplets the discounted payoff is given by

$$
\begin{aligned}
& \sum_{i=1}^{n} e^{-\int_{T_0}^{T_i} r(s)ds}\left(e^{\int_{T_{i-1}}^{T_i} r(s)ds}-(1+K\tau_i)\right) = \\
& = \sum_{i=1}^{n} e^{-\int_{T_0}^{T_{i-1}} r(s)ds}-(1+K\tau_i)e^{-\int_{T_0}^{T_i} r(s)ds}.
\end{aligned}
\tag{31}
$$

Taking expectations under the risk-neutral measure, the price at time $T_0$ is therefore

$$
\begin{aligned}
&E^{\mathbb{Q}}\left[\sum_{i=1}^{n} e^{-\int_{T_0}^{T_{i-1}} r(s)ds} - (1+K\tau_i)e^{-\int_{T_0}^{T_i} r(s)ds} | \mathcal{F}_{T_0}\right] = \\
&= \sum_{i=1}^{n} E^{\mathbb{Q}}\left[e^{-\int_{T_0}^{T_{i-1}} r(s)ds} - (1+K\tau_i)e^{-\int_{T_0}^{T_i} r(s)ds} | \mathcal{F}_{T_0}\right] = \\
&= \sum_{i=1}^{n} P(T_0, T_{i-1}) - (1+K\tau_i)P(T_0, T_i) = \\
&= \sum_{i=1}^{n} P(T_0, T_{i-1}) - P(T_0, T_i) - K\tau_i P(T_0, T_i).
\end{aligned}
\tag{32}
$$

Setting this equal to zero and solving for K gives

$$
\begin{aligned}
&\sum_{i=1}^{n} P(T_0, T_{i-1}) - P(T_0, T_i) - K\tau_i P(T_0, T_i) = 0 \\
K &\sum_{i=1}^{n} \tau_i P(T_0, T_i) = \sum_{i=1}^{n} P(T_0, T_{i-1}) - P(T_0, T_i) = P(T_0, T_0) - P(T_0, T_n) \\
K &= \frac{P(T_0, T_0) - P(T_0, T_n)}{\sum_{i=1}^{n} \tau_i P(T_0, T_i)} = \frac{1 - P(T_0, T_n)}{\sum_{i=1}^{n} \tau_i P(T_0, T_i)}.
\end{aligned}
\tag{33}
$$

This is equal to the swap rate, see section 2.9.

### 2.11.2 Pricing of Caps with Implied Volatility

A common way of quoting option prices is in terms of implied volatility rather than market prices. It is more comparable across different strikes, maturities and also currencies. However, in this analysis the model parameters are calibrated to observed market prices. In order to get the market price of a contract, the quoted implied volatility and other features of the contract have to be input into the model originally used for calculating the implied volatility. The data used in this analysis, quoted as implied volatilities, were computed using the Bachelier formula.

The Bachelier model assumes that the forward rate at time $t$, $F_t$, follows an arithmetic Brownian motion under the $T$-forward measure with volatility $\sigma$,

$$
dF_t = \sigma dW_t \tag{34}
$$

where $W_t$ is a Brownian motion under the $T$-forward measure. Below is the Bachelier formula for the price $p$ of a caplet defined on the interval $[T_1, T_2]$ given under the Bachelier model.

$$
p(0, T_1, T_2) = P(0, T_2)\tau\left[(F_0 - K)\Phi(d) + \sigma\sqrt{T_1}\phi(d)\right] \tag{35}
$$

$F_0 \coloneqq F(0; T_1, T_2)$ is the forward rate over the period between $T_1$ and $T_2$ seen from the markets perspective at $t = 0$. $\Phi$ is the standard normal cumulative distribution function, $\phi$ is the standard normal probability distribution function and $d = \frac{F_0 - K}{\sigma\sqrt{T_1}}$. $P(0, T_2)$ is the price at time 0 of a zero-coupon bond maturing at $T_2$, $\tau$ is the period between $T_1$ and $T_2$, $\sigma$ is the implied volatility and $K$ is the strike (Choi et al. (2022)). The price of the cap is given by the sum of the prices of the individual underlying caplets, all priced using the same implied volatility.

### 2.11.3 Inverting the Bachelier Formula

In order to obtain implied volatilities from the model prices, the Bachelier formula (35) is numerically inverted using the bisection method. More precisely, for a given model price $P_{mod}$ the implied volatility $\sigma_{imp}$ is defined as the value of $\sigma$ such that

$$P_{Bach}(\sigma) = P_{mod} \tag{36}$$

where $P_{Bach}$ is the Bachelier price. The function $f(\sigma)$ is then defined as

$$f(\sigma) = P_{Bach}(\sigma) - P_{mod}. \tag{37}$$

The root of the function in (37), if it exists, is found numerically using the bisection method and the implied volatility is given by the resulting solution.

### 2.11.4 Volatility Smile and Skew

In options markets implied volatility is typically not constant across strikes. When plotted against strike for a fixed maturity the graph usually shows one of two distinct patterns. One pattern, known as volatility smile, is a U-shape where implied volatility increases symmetrically when moving away from the at-the-money level and moving into the in-the-money and out-of-the-money levels. The other shape, known as volatility skew, is a shape where the implied volatility systematically increases or decreases with the strike level.

Simpler models, for example those that either assume constant volatility or contain few parameters, are generally only able to capture the ATM level and produce a flat implied volatility surface across strikes and maturities. By adding more complexity to models they are able to capture more of the smile and skew.

## 2.12 OIS Caplet Pricing Using Laplace Transform Methods

This section applies the Laplace transform and its inverse to the expression in (29) in order to obtain an alternative representation for pricing caplets (Lindström et al. (2008)). Using the following notation

$$\bar{k} = \ln\left(\frac{p_\varphi(t, T_{i-1})}{(1 + K\tau_i)p_\varphi(t, T_i)}\right) \tag{38}$$

the expression inside the inner expectation in (29) can be written as

$$\left(e^{\bar{k}} - e^{-\int_{T_{i-1}}^{T_i} X(s)ds}\right)^+. \tag{39}$$

Let $z = \gamma + i\omega$, where $\gamma, \omega \in \mathbb{R}$, be a complex number. The Laplace transform of (39) is then

$$\begin{aligned}
&\int_{-\infty}^{\infty} e^{z\bar{k}} \left(e^{\bar{k}} - e^{-\int_{T_{i-1}}^{T_i} X(s)ds}\right)^+ d\bar{k} = [U := -\int_{T_{i-1}}^{T_i} X(s)ds] = \\
&= \int_U^{\infty} \left(e^{(z+1)\bar{k}} - e^{z\bar{k}+U}\right) d\bar{k} \overset{Re(z)<-1}{=} \left[\frac{e^{(z+1)\bar{k}}}{z+1} - \frac{e^{z\bar{k}+U}}{z}\right]_U^{\infty} = \\
&= \frac{e^{(z+1)U}}{z(z+1)}.
\end{aligned} \tag{40}$$

By applying the inverse Laplace transform to the expression in (40), the expected value of the expression in (39) under the risk-neutral measure conditional on the filtration at time $T_{i-1}$ is

$$E^{\mathbb{Q}}\left[\left(e^{\bar{k}} - e^{U}\right)^{+} | \mathcal{F}_{T_{i-1}}\right] = \frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{e^{-z\bar{k}}}{z(z+1)} E^{\mathbb{Q}}\left[e^{-(z+1)\int_{T_{i-1}}^{T_i} X(s)ds} | \mathcal{F}_{T_{i-1}}\right] d\omega. \tag{41}$$

Now, assuming that X(t) follows a process belonging to the class of affine term structure models, the conditional expectation in (41) admits a closed-form expression. Using this, the expression in (41) simplifies to

$$\frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{e^{-z\bar{k}}}{z(z+1)} e^{\tilde{A}(T_{i-1},T_i)+\tilde{B}(T_{i-1},T_i)X(T_{i-1})} d\omega. \tag{42}$$

Substituting this into (29) yields

$$\begin{aligned} &E^{\mathbb{Q}}\left[(1+K\tau_i)p_\varphi(t,T_i)e^{-\int_t^{T_{i-1}} X(s)ds}\frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{e^{-z\bar{k}}}{z(z+1)} e^{\tilde{A}(T_{i-1},T_i)+\tilde{B}(T_{i-1},T_i)X(T_{i-1})} d\omega | \mathcal{F}_t\right] = \\ &= (1+K\tau_i)p_\varphi(t,T_i)\frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{e^{-z\bar{k}}}{z(z+1)} d\omega E^{\mathbb{Q}}\left[e^{-\int_t^{T_{i-1}} X(s)ds + \tilde{A}(T_{i-1},T_i)+\tilde{B}(T_{i-1},T_i)X(T_{i-1})} | \mathcal{F}_t\right] = \\ &= (1+K\tau_i)p_\varphi(t,T_i)\frac{1}{2\pi}\int_{-\infty}^{\infty} \frac{e^{-z\bar{k}}}{z(z+1)} e^{\bar{A}(t,T_{i-1})+\bar{B}(t,T_{i-1})X(t)} d\omega. \end{aligned} \tag{43}$$

## 2.13 The CIR++ Model

For the analysis, the CIR++ model is used. CIR++ is a one-factor short-rate model with deterministic shift, as defined in section 2.6, where the stochastic process $X_t$ follows a CIR process. The dynamics of the CIR process under the risk-neutral measure $\mathbb{Q}$ are given below.

$$dX_t = \kappa(\theta - X_t)dt + \sigma\sqrt{X_t}dW_t \tag{44}$$

Parameter $\kappa$ is the mean reversion speed, $\theta$ is the long-term mean, $\sigma$ is the volatility parameter and $W_t$ is a Brownian motion under $\mathbb{Q}$.

From equation (19) it can be seen that the CIR-process belongs to the class of affine term-structure models. By setting

$$\alpha(t) = -\kappa, \beta(t) = \kappa\theta, \gamma(t) = \sigma^2, \delta(t) = 0 \tag{45}$$

in (19) results in equation (44). The equations for functions $A$ and $B$ in (21) become

$$\begin{aligned} B_t'(t,u) &= \kappa B(t,u) + \frac{1}{2}\sigma^2 B^2(t,u) - 1,\ B(u,u) = 0, \\ A_t'(t,u) &= \kappa\theta B(t,u),\ A(u,u) = 0. \end{aligned} \tag{46}$$

These equations admit closed form solutions, see appendix A.1.

The pricing formula in (43) requires the value of $X_t$ at the valuation time $t$. Since $X_t$ is not directly observable, the continuously compounded zero-coupon rate at time $t$ corresponding to the shortest observed maturity is used as a proxy.

One important advantage of the CIR++ model is the exogenous nature of it due to $\varphi$. While models like Vasicek and CIR produce an endogenous term structure of interest rates, meaning that these models create their own term structure that needs to be fitted, models like CIR++ where the term structure is an input can use the other parameters to calibrate to other market instruments.

When prices are computed of interest rate contingent claims, they are calculated under the risk-neutral measure. Therefore, when calibrating the model parameters using derivative prices, like cap prices, this must be done under the risk-neutral dynamics in accordance with (44).

## 2.14 Calibration

This section describes the calibration of the CIR++ model to market cap prices. The objective is to determine the model parameters such that modeled cap prices are consistent with the observed market prices. The parameters $\kappa$ and $\theta$ were fixed at the values $\kappa = 1$ and $\theta = 0.03$, while $\sigma$ was modeled as piecewise constant. Alternative values of $\kappa$ and $\theta$ were also examined, but as their effect on the calibration and pricing results was limited, both parameters were kept fixed at the selected values.

### 2.14.1 Fitting the Yield Curve

The CIR++ model consists of two parts, see (16). As a result of how the deterministic term is defined, the model is perfectly calibrated to the zero coupon curve at the calibration date, see (18). By choosing the calibration date to coincide with the valuation date therefore ensures that the model perfectly reproduces the initial term structure of interest rates.

### 2.14.2 Calibration of the Volatility Parameters

The volatility parameter $\sigma$ is modeled as piecewise constant, i.e.

$$\sigma(t) = \sum_{i=1}^{N} \sigma_i \mathbf{1}_{[T_{i-1},T_i)}(t). \tag{47}$$

The time points $T_i$ are chosen to coincide with those defining the tenor structure of the caps introduced in section 2.11. The first volatility parameter $\sigma_1$ is defined on the interval $[T_0, T_1]$, the second parameter $\sigma_2$ on $[T_1, T_2]$ and so forth. For a full specification of how the volatility parameters and the tenor structure are defined, see table (6) in Section A.3.

Assuming a piecewise constant $\sigma(t)$ allows for sequential calibration. First, $\sigma_1$ is calibrated using caps with a maturity of one year. Next, $\sigma_2$ is calibrated using caps with a maturity of two years while keeping $\sigma_1$ fixed. This is repeated iteratively for all volatility parameters. Since calibration is performed using ATM caps it is not possible to calibrate for example $\sigma_2$ using only the additional caplets contained in caps with a maturity of two years compared to caps with a maturity of one year. This is due to the fact that ATM caps have different strikes.

Since $\sigma(t)$ is piecewise constant, this has to be taken into consideration when discounting. Discounting is performed iteratively, starting from the end and proceeding backwards. Assuming a tenor structure consisting of $n$ maturities the procedure is as follows. First, discounting between $T_n$ and $T_{n-1}$ is performed using expression (20). The functions $A$ and $B$ in the zero-coupon bond price (20) satisfy equations (21) with corresponding boundary conditions and volatility parameter. When discounting over the interval $[T_{n-2}, T_{n-1}]$, functions $A$ and $B$ also satisfy equations (21) but with non-zero boundary conditions. The boundary values of $A$ and $B$ are given by the values at the endpoint of the preceding interval. Discounting is then performed in the same way for the all the remaining intervals until the beginning. The intervals follow the same structure as the volatility parameters, because $A$ and $B$ are dependent on them.

In general, the procedure above could be carried out for each caplet rather than the intervals corresponding to the tenor structure and the parameters. However, since the parameters are constant over several caplets it is more efficient to perform discounting based on the intervals. The pricing routine handles this by partitioning the discounting interval into as few subintervals as possible.

The calibration of each parameter is performed in Matlab using the built-in optimization function lsqnonlin, which is a nonlinear least squares method. The expression for the caplet price in (43) is solved numerically using a Gauss-Laguerre quadrature method.

Allowing $\sigma(t)$ to be piecewise constant increases the flexibility of the model while preserving the analytical tractability. By modeling $\sigma(t)$ as piecewise constant, the functions $A$ and $B$ can be solved analytically on each interval.

# 3 Data Description

This section starts with an overview of time conventions and how time to maturity is measured, it also specifies the convention used in the analysis. It then goes on to present the datasets that are used. All market data in this study was sourced directly from the Bloomberg Terminal. This data was also partly used in Stoltz and Sönne (2025).

## 3.1 Time Convention and Day-Count Methodology

When referring to points in time such as $t$ and $T$, the notation has different meanings. Sometimes they refer to real numbers that are measured in relation to an instant that has been chosen as origin $t = 0$, while other times they are referred to as dates. In practice, when working with data, time differences such as time to maturity $\tau$ are typically expressed as year fractions. The general idea behind the year fraction is to take the number of days between two dates and divide by the number of days in a year. There are different ways how this is specified and they are called day-count conventions. One example of such day-count conventions is the Actual/360, which is the one used for the data in this study. Under this convention, a year is assumed to consist of 360 days. Let $D_1$ and $D_2$ denote two dates, and let $D_2 - D_1$ denote the number of days between them, where $D_1$ is included and $D_2$ is excluded. The year fraction is then defined as (Brigo and Mercurio (2006))

$$\frac{D_2 - D_1}{360}. \tag{48}$$

## 3.2 Data Categories

### 3.2.1 SOFR Overnight Indexed Swap (OIS) Rates

The SOFR OIS rates were obtained in the form of bid and ask quotes derived from observed market transactions. The contracts follow an Actual/360 day-count convention with annual payment frequency and uses a backward rolling convention. A settlement lag of two business days is used, and all rates reference the SOFR index.

The sample period extends from January 1, 2021 to February 7, 2025 and covers a wide range of maturities:

- **Short maturities:** 1D, 2D, 3D, 1W, 2W, 3W
- **Monthly maturities:** 1M to 11M; 13M to 23M; as well as 27M, 30M, 42M, and 54M
- **Yearly maturities:** 1Y to 12Y at annual intervals
- **Long-dated maturities:** 15Y, 20Y, 25Y, 30Y, 35Y, 40Y, 45Y, and 50Y

### 3.2.2 SOFR Cap Normal Volatility Data

Normal volatilities for SOFR caps were obtained from Bloomberg's VCUB. The quoted volatilities are expressed in basis points (bps) and are provided within the Bachelier (normal) framework. The underlying reference rate is the compounded SOFR index.

The volatility surface is constructed from mid-market quotes. For maturity–strike combinations that are not directly traded, implied volatilities are generated through interpolation. This interpolation is performed using a SABR model in the normal volatility space.

The contracts follow the Actual/360 day-count convention and apply a backward-looking compounding methodology for the overnight rate. It assumes quarterly payments and a settlement lag of two business days.

The sample period spans from July 22, 2021 to February 3, 2025. The tenor structure includes maturities that range from 1 to 10 years in annual steps, with additional maturities at 12, 15, 20, and 30 years.

Strike levels are quoted in absolute terms and range from $-200$ bp to 700 bp. The grid contains the following strikes: $-200$, $-100$, $-50$, $-25$, 0, and ATM, followed by increments of 25 bp from 25 bp to 300 bp, and subsequently 50 bp increments up to 700 bp.

## 3.3 Preprocessing

The mid-price was calculated from the bid and ask prices as

$$Mid = \frac{Bid + Ask}{2}. \tag{49}$$

The normal volatilities were converted into cap prices for a notional amount of USD 1 using expression (35). The resulting prices are expressed in U.S. dollars. Since the datasets were observed on different dates, they were processed to ensure that all data were aligned to a common set of dates.

# 4 Results

This section presents the results of the empirical analysis. Two different sample periods are considered, corresponding to two distinct market regimes characterized by different yield curve shapes. By examining periods with different yield-curve shapes, the analysis evaluates how the model performs under different market conditions. The performed tests focus on the model's in-sample fit, its out-of-sample pricing performance and whether the inclusion of additional calibration data improves pricing accuracy.

## 4.1 Market Environments

In order to examine how the model performs under different market conditions, two sample periods are considered in the analysis, see Table 1. Both periods consist of 40 days and are characterized by different yield-curve shapes, one normal and one inverted. In addition, one representative date is selected from each period based on the shape of the yield curve. These dates are used both to illustrate the characteristic yield-curve shape of each period and to study the model's performance more closely under the two market conditions.

| Period | Start date | End date | Representative date |
|---|---|---|---|
| Period 1 | 2021-09-15 | 2021-11-17 | 2021-09-20 |
| Period 2 | 2024-06-03 | 2024-07-29 | 2024-06-10 |

Table 1: Sample periods and representative dates used in the empirical analysis.

In the first period, the yield curve is broadly characterized by a normal shape, with a generally upward-sloping pattern over the relevant maturities. The yield curve for the representative date in Period 1 is shown in Figure 1.

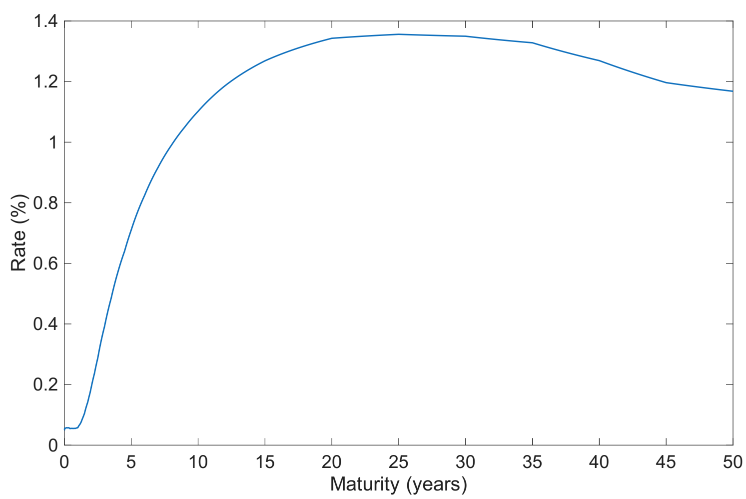


Figure 1: Yield curve for the representative date in Period 1.

In the second period, the yield curve generally exhibits a typical inverted shape. The yield curve for the representative date in Period 2 is shown in Figure 2.

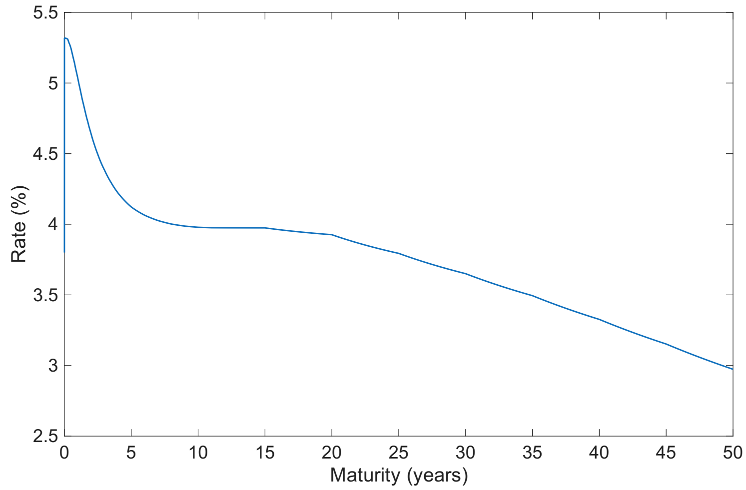


Figure 2: Yield curve for the representative date in Period 2.

## 4.2 In-Sample Fit

To assess the in-sample performance of the model, the volatility parameters were calibrated for each day in the two sample periods using ATM contracts. The calibrated parameters were then used to price the corresponding contracts across all maturities. For each maturity, the root mean square error (RMSE) was computed across the 40 days in each period, providing a quantitative measure of the fit between the model prices and the corresponding market prices. In addition, for the representative day in each period, the market prices and model prices are presented, as well as the corresponding residuals.

Figures 3a and 3b below present the market and model prices plotted against maturity and the corresponding residuals, respectively, for the representative date in Period 1. The calibrated volatility parameters are shown in Figure 15a and the ATM strikes are listed in Table 7.

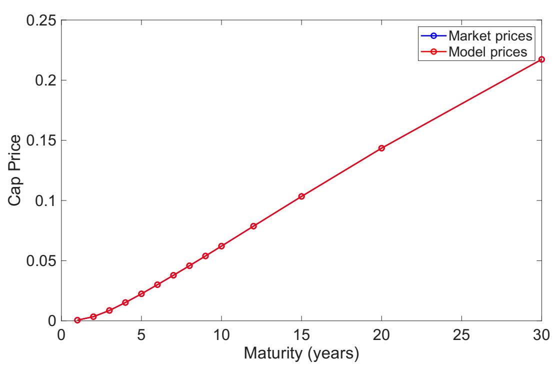


(a) Market and model prices.

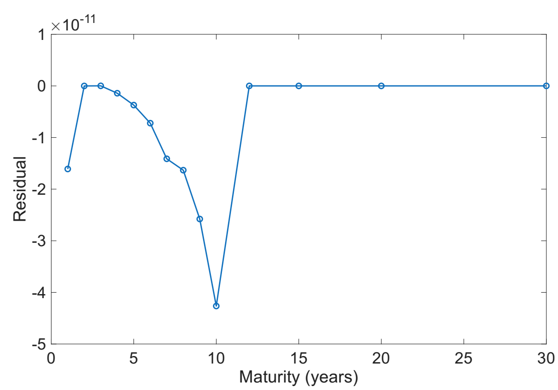


(b) Residuals.

Figure 3: In-sample market and model cap prices across maturities and corresponding residuals for the representative date in Period 1.

For the representative date in Period 2, the corresponding figures are presented in Figures 4a and 4b. The calibrated volatility parameters are shown in Figure 15b and the ATM strikes are listed in Table 7.

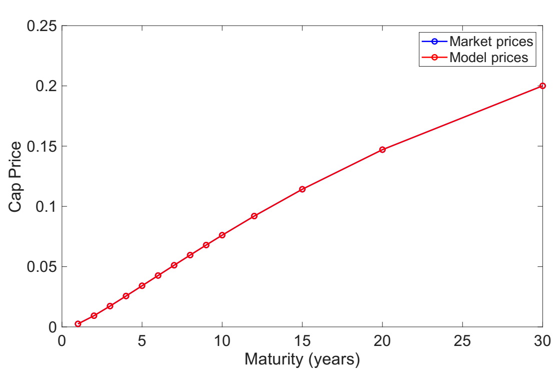


(a) Market and model prices.

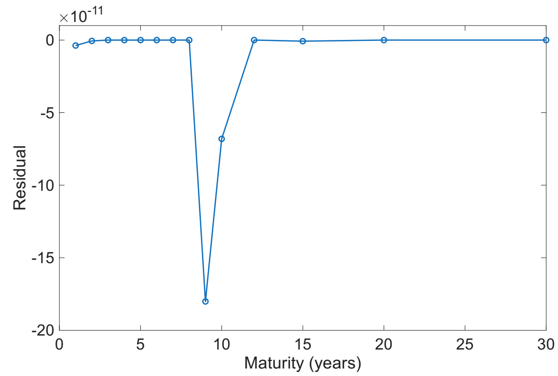


(b) Residuals.

Figure 4: In-sample market and model cap prices across maturities and corresponding residuals for the representative date in Period 2.

Figures 5a and 5b show the RMSE values across maturities for Periods 1 and 2, respectively.

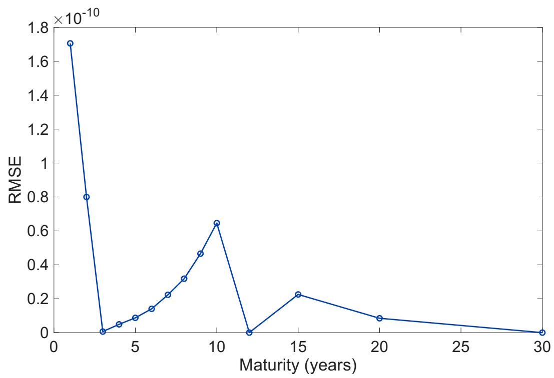


(a) In-sample RMSE across maturities for Period 1.

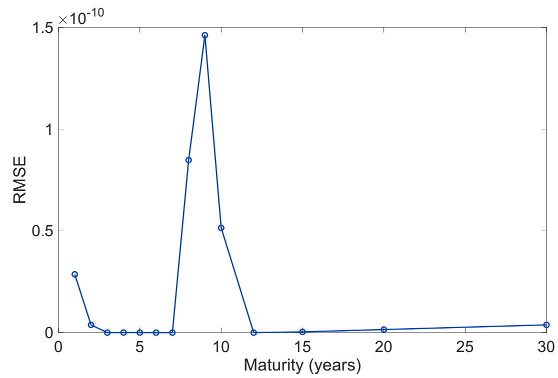


(b) In-sample RMSE across maturities for Period 2.

Figure 5: In-sample RMSE across maturities for both sample periods.

## 4.3 Cap Price Surface

To assess the model's ability to price contracts out of sample across strikes, the parameters were calibrated to different sets of strikes across all maturities and subsequently used to price contracts at additional strike levels. The analysis was performed on the representative dates in Period 1 and 2.

### 4.3.1 Period 1

Figure 6 presents the market and model prices for the representative date in Period 1 when the model is calibrated to ATM strikes. In addition, prices corresponding to the ATM contracts are shown. The strikes used for the pricing range from 0 to 700 bps and the ATM strikes lie roughly between 5 and 135 bps. In Period 1, data are missing for the strikes $K = 250$ and $K = 350$, as reflected in the figure.

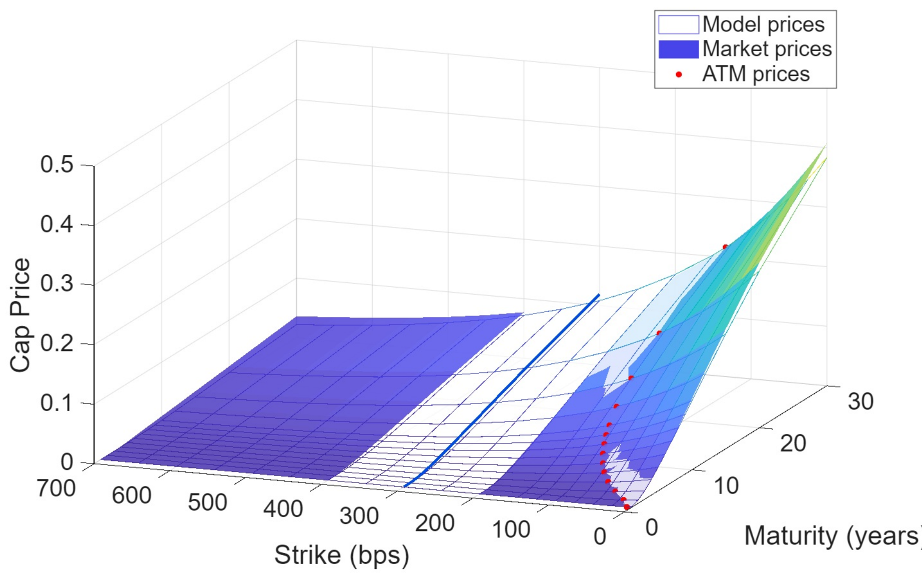


Figure 6: Market and model cap prices across strikes and maturities for the representative date in Period 1, when the model is calibrated to ATM strikes.

Four different sets of strikes are considered for the calibration. The precise composition of these sets is given in Table 2.

| Strike set | Included strikes (bps) |
|---|---|
| Set 1 | ATM (approximately 5 to 135) |
| Set 2 | ATM, 75, 200 |
| Set 3 | ATM, 25, 700 |
| Set 4 | ATM, 300, 400 |

Table 2: Strike sets used in the calibration for the representative date in Period 1.

Figure 7 shows RMSE across maturities and strikes for the different calibration sets.

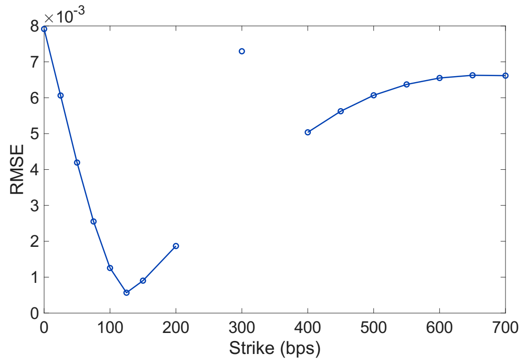


(a) RMSE across strikes, strike set 1.

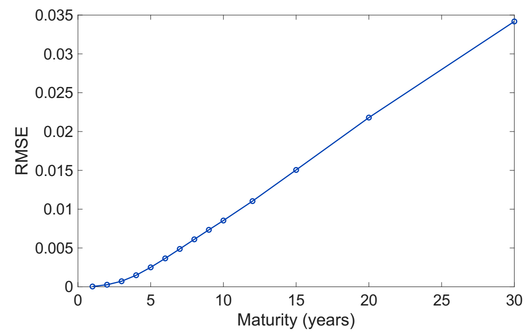


(b) RMSE across maturities, strike set 1.

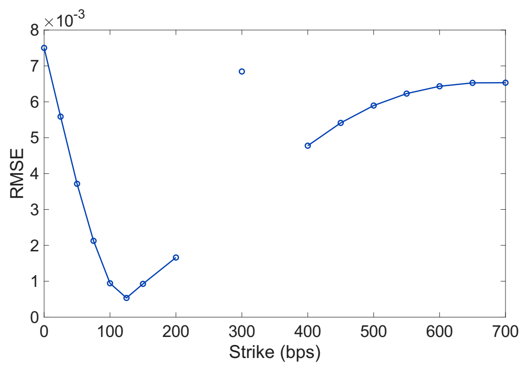


(c) RMSE across strikes, strike set 2.

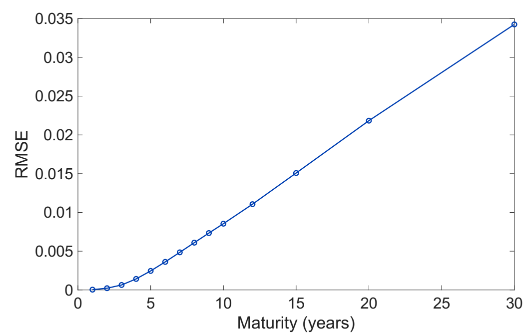


(d) RMSE across maturities, strike set 2.

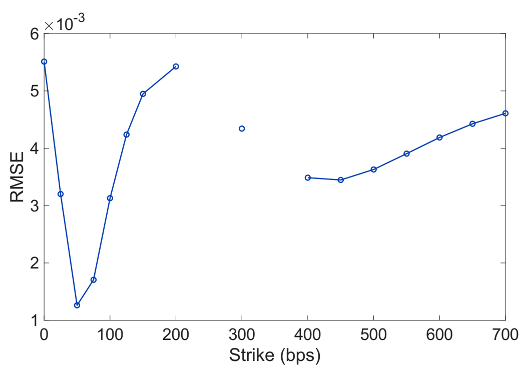


(e) RMSE across strikes, strike set 3.

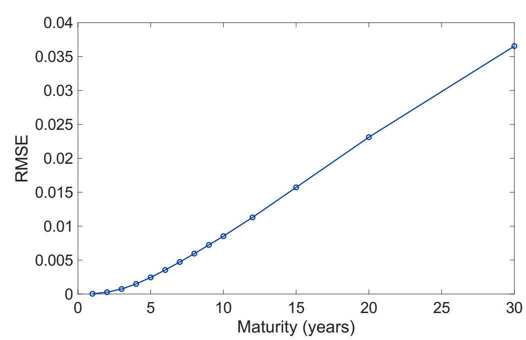


(f) RMSE across maturities, strike set 3.

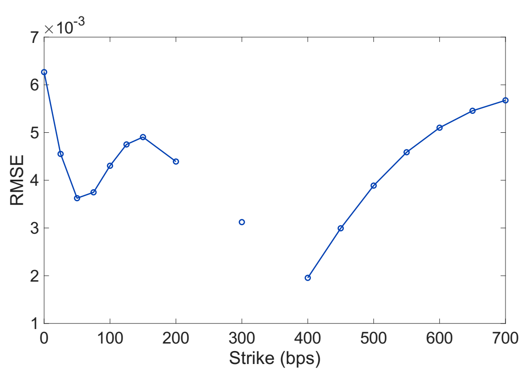


(g) RMSE across strikes, strike set 4.

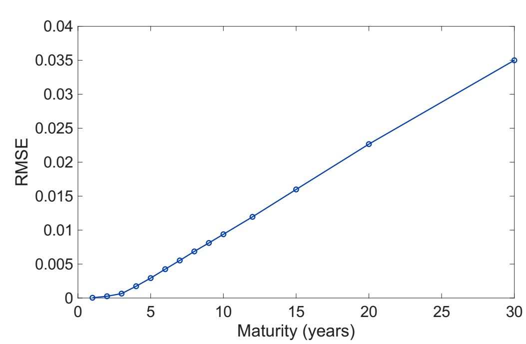


(h) RMSE across maturities, strike set 4.

Figure 7: RMSE across strikes and maturities for the representative date in Period 1 under the different calibration strike sets.

Table 3 lists the total RMSE values between model and market price surfaces under the different calibration strike sets for the representative date in Period 1.

| Strike set | RMSE |
|---|---|
| Set 1 | 0.0125 |
| Set 2 | 0.0126 |
| Set 3 | 0.0132 |
| Set 4 | 0.0131 |

Table 3: Total RMSE between model and market price surfaces for the different calibration strike sets for the representative date in Period 1.

### 4.3.2 Period 2

Figure 8 presents the market and model prices for the representative date in Period 2 when the model is calibrated to ATM strikes. In addition, prices corresponding to the ATM contracts are shown. The strikes used for the pricing ranges between $K = 25$ and $K = 700$.

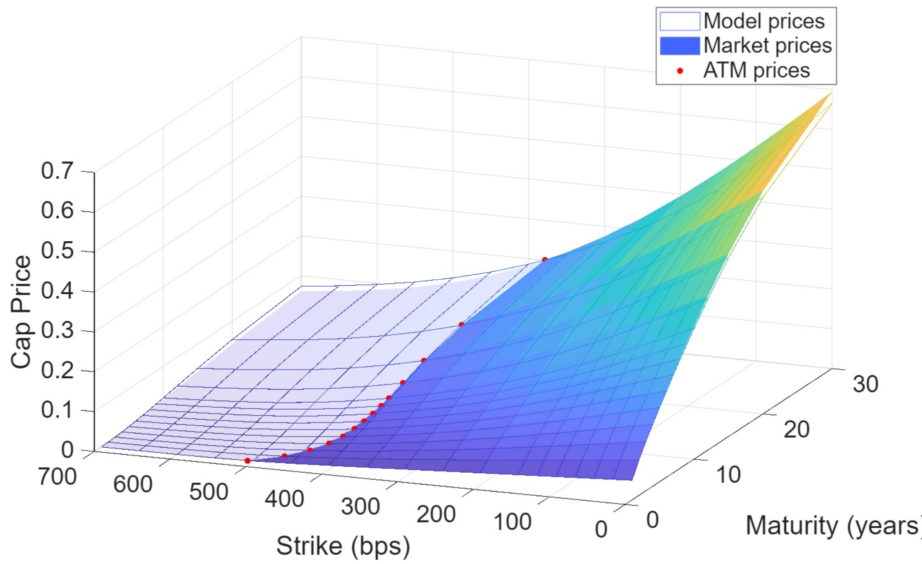


Figure 8: Market and model cap prices across strikes and maturities for the representative date in Period 2, when the model is calibrated to ATM strikes.

The composition of the four different sets of strikes used in the calibration is given in Table 4.

| Strike set | Included strikes (bps) |
|---|---|
| Set 1 | ATM (approximately 380 to 510) |
| Set 2 | ATM, 300, 550 |
| Set 3 | ATM, 25, 700 |
| Set 4 | ATM, 25, 100 |

Table 4: Strike sets used in the calibration for the representative date in Period 2.

Figure 9 shows RMSE across maturities and strikes for the different calibration sets.

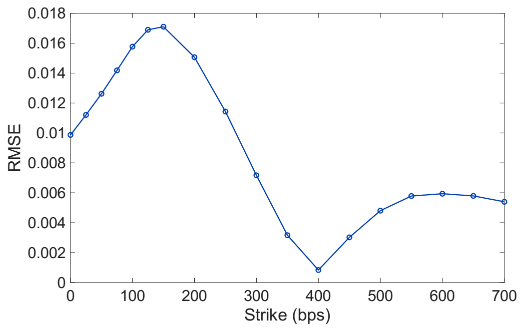


(a) RMSE across strikes, strike set 1.

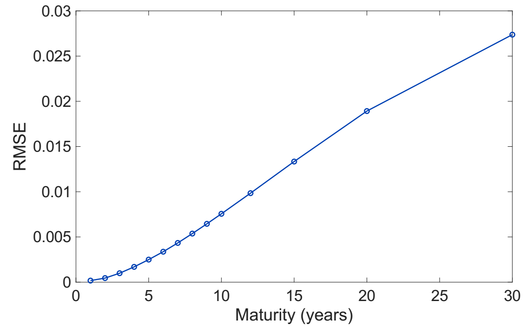


(b) RMSE across maturities, strike set 1.

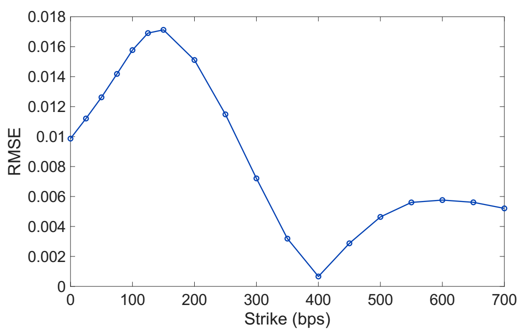


(c) RMSE across strikes, strike set 2.

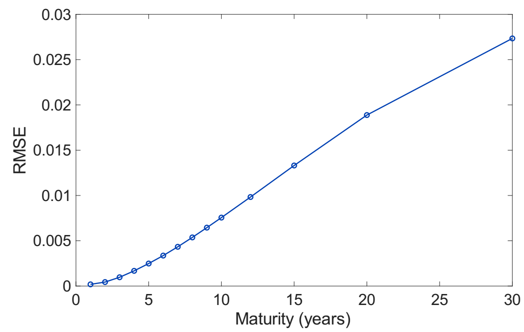


(d) RMSE across maturities, strike set 2.

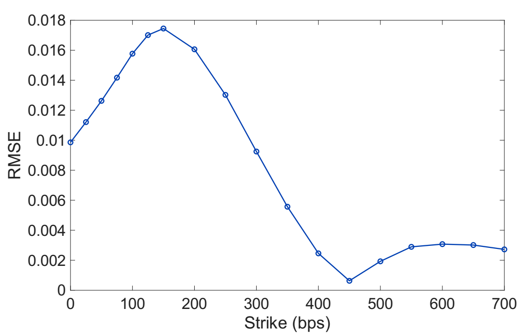


(e) RMSE across strikes, strike set 3.

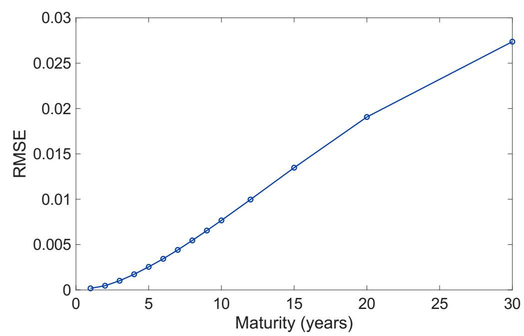


(f) RMSE across maturities, strike set 3.

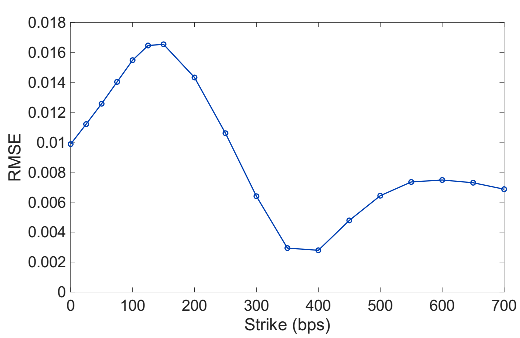


(g) RMSE across strikes, strike set 4.

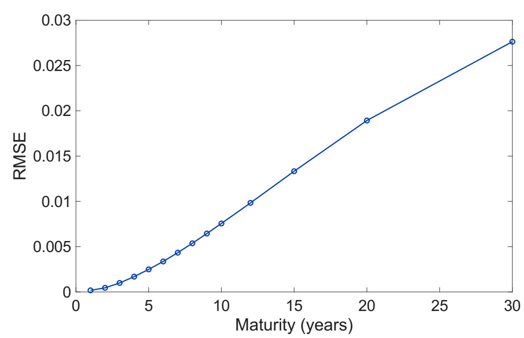


(h) RMSE across maturities, strike set 4.

Figure 9: RMSE across strikes and maturities for the representative date in Period 2 under the different different calibration strike sets.

Table 5 lists the total RMSE values between model and market price surfaces under the different calibration strike sets for the representative date in Period 2.

| Strike set | RMSE |
|---|---|
| Set 1 | 0.0105 |
| Set 2 | 0.0105 |
| Set 3 | 0.0106 |
| Set 4 | 0.0106 |

Table 5: Total RMSE between model and market price surfaces for the different calibration strike sets for the representative date in Period 2.

## 4.4 Effect of the Number of Calibration Days on Out-of-Sample Pricing

Next, an additional test is performed to assess the model's ability to price contracts outside the calibration set. The purpose is both to evaluate the model's predictive performance on dates not included in the calibration and to examine whether using a longer calibration window improves pricing accuracy. The model is calibrated using ATM contracts from either one day or five consecutive days. In the latter case, the calibration window consists of Monday to Friday within a given week, whereas the one-day calibration is based on the Friday of that same week. Then, using the calibrated parameters, the corresponding contracts for the following Monday are priced. This procedure is repeated over a rolling window throughout the 40-day sample in both periods. Figure 10 shows the RMSE values by maturity for both periods under the one-day and five-day calibration schemes. The corresponding calibrated volatility parameters are shown in Figure 16.

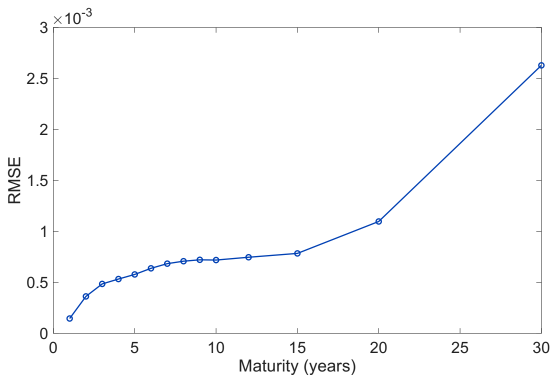


(a) Five-day calibration for the first period.

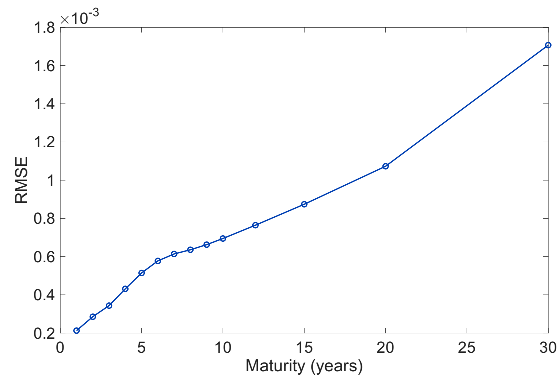


(b) One-day calibration for the first period.

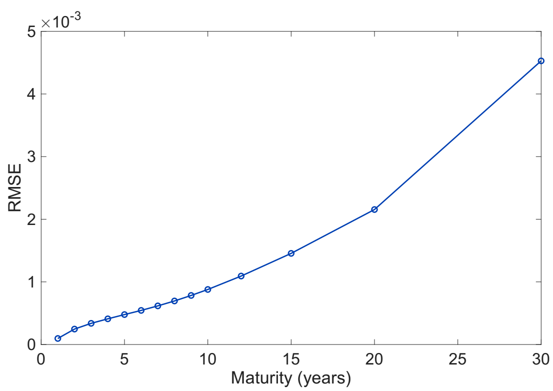


(c) Five-day calibration for the second period.

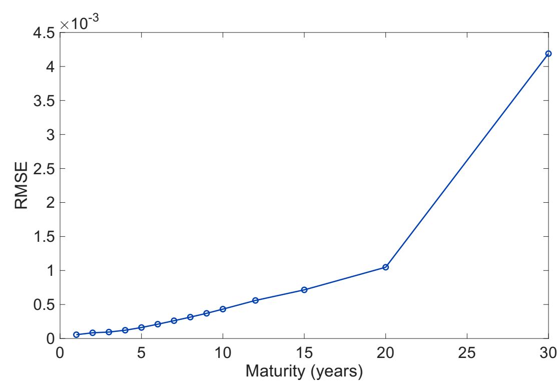


(d) One-day calibration for the second period.

Figure 10: RMSE by maturity for both periods under the one-day and five-day calibration schemes.

For the representative dates in Periods 1 and 2, the market and model prices are plotted for both calibration schemes, and the corresponding residuals are presented separately. Figures 11a and 11b show the prices and residuals, respectively, for the representative date in Period 1.

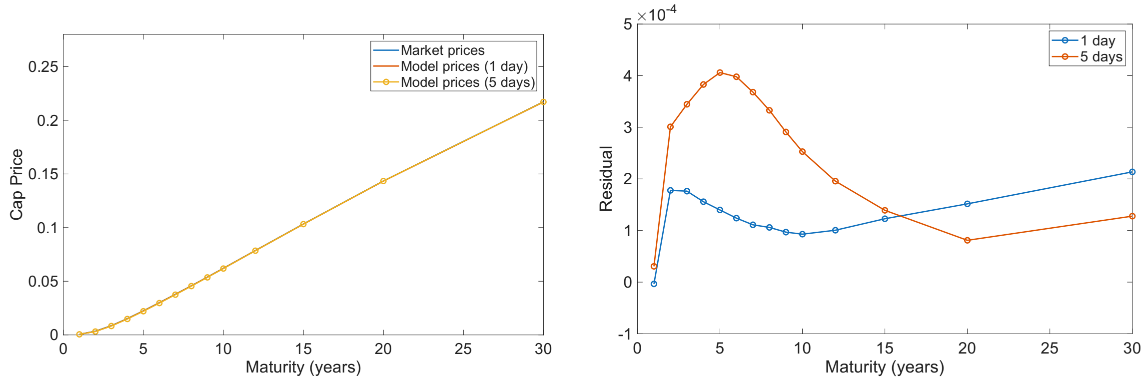


(a) Market and model cap prices across maturities.

(b) Residuals across maturities.

Figure 11: Market and model cap prices and corresponding residuals for the representative date in Period 1 under the one-day and five-day calibration schemes.

For the representative date in Period 2, the corresponding results are shown in Figures 12a and 12b.

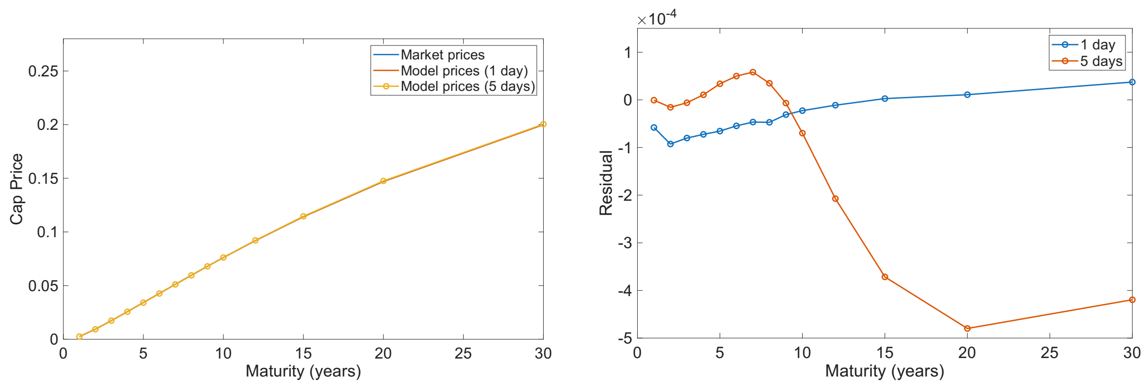


(a) Market and model cap prices across maturities.

(b) Residuals across maturities.

Figure 12: Market and model cap prices and corresponding residuals for the representative date in Period 2 under the one-day and five-day calibration schemes.

## 4.5 Absolute and Relative Pricing Errors

This section presents figures of absolute and relative pricing errors in order to assess whether these two error measures provide a consistent picture across maturities and strikes. The results for the representative date in Period 1 are presented in Figure 13.

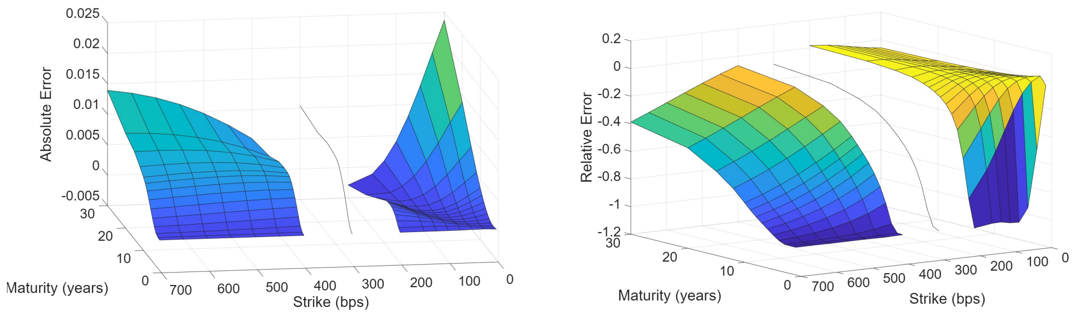


(a) Absolute errors.

(b) Relative errors.

Figure 13: Absolute and relative pricing errors between market and model cap prices for the representative date in Period 1 under calibration to ATM strikes.

The results for the representative date in Period 2 are presented in Figure 14.

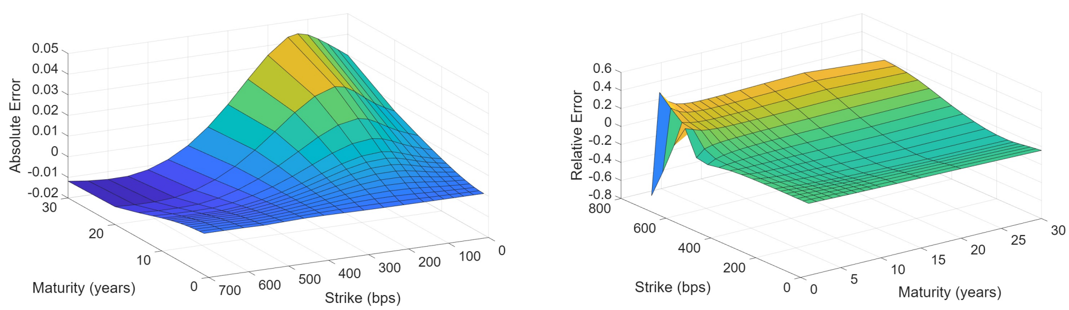


(a) Absolute errors.

(b) Relative errors.

Figure 14: Absolute and relative pricing errors between market and model cap prices for the representative date in Period 2 under calibration to ATM strikes.

# 5 Discussion and Conclusion

## 5.1 In-Sample Fit

Overall, the results suggest that the model is able to reproduce the market cap prices well, although the fit varies somewhat between maturities, strikes and market conditions. As seen from Figures 5a and 5b the RMSE values are generally low across maturities in both periods, although some variation can be observed. In particular, the RMSE values are somewhat elevated at the short end, especially for the one- and two-year maturities. One reason for this could be the sensitivity to the approximation of the initial interest rate. At the intermediate maturities the pattern differs between the two periods. In Period 1, the RMSE increases gradually up to a peak at the ten-year maturity, followed by a rapid decline. In Period 2, the RMSE values remain generally low up to a peak at the nine-year maturity, also followed by a rapid decline. At the long end in Period 1, a smaller elevation is observed at the fifteen-year maturity, after which the RMSE decreases gradually for longer maturities. In Period 2, on the other hand, the RMSE increases gradually with maturity.

In the model, $\kappa$ and $\theta$ are kept fixed, leaving 14 parameters to be calibrated, one volatility parameter for each maturity. Since there exists a unique contract for each maturity, the number of parameters matches the number of contracts and this likely contributes to the strong in-sample fit.

## 5.2 Cap Price Surface

From Figures 6 and 8 the model appears to capture the overall shape of the cap price surface well. Although the calibration is based only on ATM strikes, the resulting parameter estimates still allow the model to reproduce prices across a broad range of non-ATM strikes with reasonable accuracy. The largest deviations are observed for the longer maturities. For both periods, the RMSE values across maturities, shown in Figures 7 and 9, support this observation, as the RMSE increases consistently with maturity. The deviation is smallest around the ATM strikes and increases as the strike moves away from the ATM levels. This can also be observed from the the figures showing RMSE across strikes in Figures 7 and 9. They generally all attain the minimum values around the ATM levels.

In Period 1, Figures 7a and 7c show that the RMSE across strikes exhibits a similar shape for calibration Sets 1 and 2, although the values are somewhat lower for Set 2 in the lower and middle parts of the strike range. This is reasonable since Set 2 contains additional strikes at $K = 75$ and $K = 200$ bps. For the third set, the shape of the RMSE across maturities in Figure 7e shifts slightly and the RMSE values get somewhat lower for the lower and higher strikes consistent with the additional strikes in Set 3. For the forth set, the shape is also different. The RMSE values are slightly lower for the medium strikes, which is reasonable since the additional strikes at $K = 300$ and $K = 400$ are in that area. There is also an increase for the lower strikes. For the different calibrations sets, the effect on the RMSE values across maturities is generally small, as indicated by the maturity subfigures in Figure 7.

In Period 2, for the second calibration set, which includes additional strikes a little below and above the ATM levels, Figure 9c shows that there is little effect on the RMSE values across strikes. For the third set, which includes strikes at the lower and higher ends of the strike range, Figure 9e shows that the inclusion of these strikes improves pricing for the higher strikes while the lower end is not affected much. For the fourth set, containing two additional strikes at the lower end, Figure 9g shows that the effect on the RMSE values for the lower strikes $K = 25, 50, 75$ is small. While, for the strikes between $K = 100$ and $K = 350$ the effect is greater. For the higher strikes, $K = 400$ and above, the RMSE values come close to those in Figures 9a and 9c, where the calibration strikes are all around the ATM levels. As shown in the maturity subfigures in Figure 9, the different calibrations sets appear to have little effect on the RMSE values.

The overall effect from using different sets of calibration strikes seems to be small, indicated by Tables 3 and 5 for Period 1 and 2 respectively. For Period 1, the RMSE values become slightly worse by including additional strikes in the calibration, whereas for Period 2 they generally remain unchanged. When calibration is performed using other strikes in addition to ATM strikes, it is possible for $\kappa$ and $\theta$ to vary with maturity, in the same way as the volatility parameter, while still keeping the number of contracts equal to the number of parameters. Figure 15, which shows the calibrated volatility parameters, suggests that the values do not vary substantially across maturities and the different calibrations sets. The pattern across maturities is generally similar for all calibrations in both periods, although the general level differs between the two periods. One notable feature is a slight oscillation at the intermediate maturities in Figure 15g. Otherwise, the parameters appear to be rather stable.

Another point to note is that, for certain contracts, the Bachelier formula could not be inverted to obtain the implied volatility. For this reason the results in this section are presented in terms of prices rather than implied volatilities.

## 5.3 Effect of the Number of Calibration Days on Out-of-Sample Pricing

Figures 10a to 10d show that when pricing contracts on dates outside of the calibration set the model appears to perform reasonably well, although its in-sample pricing performance is considerably better. The RMSE values increase from the order of $10^{-11}$ when pricing and calibrating on the same day to the order of $10^{-3}$ when pricing on a different day. Across all calibrations there is a pattern where RMSE is increasing with maturity, which may reflect the greater uncertainty associated with interest rates over longer time horizons. For both periods, when calibrating using five days of data the corresponding RMSE-values are higher than the ones obtained when calibrating using a single day. While calibration to a single day may increase the risk of overfitting, the empirical results seem to suggest that time variation in the parameters and market conditions is the dominant effect. Therefore, calibrating to one day appears to capture the current market conditions better than calibrating to multiple days. The results in Figures 11 and 12 further support the conclusion that the model is able to reproduce market prices reasonably well in both periods. In Figure 16 the values of the volatility parameters are shown for both calibrations in both periods. The parameter values are similar across the two calibrations in both periods and the overall pattern and the levels are consistent with those observed in Figure 15. This supports the observation that the volatility parameters appear to be stable.

## 5.4 Absolute and Relative Pricing Errors

For the representative date in period 1, Figure 13a shows that the absolute errors generally seem to increase with maturity. In terms of strike, the absolute errors appear to increase with strike, except at the very short end, where the errors are typically larger. For the shortest maturities, the errors vary only little with strike. Figure 13b shows that the relative errors tend to decrease with maturity and increase with strike.

For the representative date in Period 2, Figure 14a suggests that absolute error generally increases with maturity for low and medium strikes, whereas for higher strikes they are less variable or even slightly decreasing. In terms of strike, the absolute errors are typically largest for the lowest strikes. They then decrease towards the medium strikes, where the errors are lowest, before increasing again for higher strikes. For the shortest maturities, the absolute errors does not vary much with strike. Figure 14b shows that the relative errors seem to not vary much with maturity, or slightly decrease, but increase with strike.

The observation that absolute errors tend to increase with maturity, while relative errors tend to decrease with maturity, suggests that the model generally prices contracts with longer maturities better in relative terms. It also indicates that the higher absolute errors observed at longer maturities are primarily driven by the higher cap prices associated with those maturities.

## 5.5 Conclusion

The thesis has aimed to analyze the ability of the proposed short rate model, a CIR++ model with piecewise constant volatility parameter, to price OIS caps. Particular emphasis has been on in-sample fit, out-of-sample pricing performance and the effect of different calibration sets and market conditions. Taken together, the results from the analysis suggest that the model is able to reproduce contracts reasonably well, although the fit varies across maturities, strikes, and market conditions.

The in-sample results show that the model generally provides a good fit to market cap prices across maturities in both sample periods. While there is some variation for the different maturities, the RMSE values generally remain low. This indicates that the model is able to capture the term structure of the cap prices reasonably well. However, part of the strong fit is likely attributable to the calibration setup, since the number of volatility parameters matches the number of calibration contracts.

The out-of-sample analysis show that the model is also able to price caps on dates that is not included in the calibration set with reasonable accuracy, although the performance is considerably less accurate than in sample. Comparing the results from calibrating to one day or five days of cap prices, the one-day calibration tends to performs better, indicating that the time variation in parameters and market conditions makes it better to calibrate to a single recent day rather than averaging over several. The results also suggest that the model captures the surface structure across strikes and maturities with reasonable accuracy. When other strikes in addition to the ATM strikes are included in the calibration, the RMSE across strikes varies with the setup, but the overall performance mostly remains the same or slightly worse. This may indicate that it is sufficient to calibrate using only ATM strikes.

The comparison of absolute and relative errors provides additional insight into the performance across maturities. The absolute errors generally increase with maturity while relative errors tend to decrease with maturity. This observation suggests that the model in relative terms prices caps with longer maturities better and that the higher prices associated with contracts with longer maturities contributes to the bigger absolute errors.

The volatility parameters appear to be stable across maturities and over time. They exhibit a similar pattern across all calibration sets used in section 4.3, where they increase with maturity in a relatively uniform manner. The increments are slightly higher for the short end of the maturity range but otherwise they are generally similar. The same overall pattern is also observed when the parameters are calibrated out of sample using both one day and five days of data.

Overall, the results indicate that the model provides an analytic tractable and reasonably accurate framework for pricing caps in terms of prices rather than implied volatility. The model is able to price in sample with good accuracy, although part of this strong fit is likely due to the fact that the number of volatility parameters matches the number of contracts used in the calibration. The out-of-sample analysis show that the model is able to capture the market cap prices relatively well both , although it is considerably less accurate than in-sample. The fit varies across the price surface and is dependent on market conditions and calibrations. When pricing contracts on dates outside the calibration set, the results suggest that calibrating to a single recent date yields better performance than calibrating to several dates. The calibrated volatility parameters remain relatively stable throughout the different tests, both across maturities and over time.

# 6 Further Research

Since the Bachelier formula could not always be inverted to obtain the implied volatility for all contracts, it would be relevant to examine this issue further. It would also be interesting to study the model's performance more thoroughly in implied-volatility terms and to consider calibrating using implied volatilities.

A possible extension of the present analysis would also be to calibrate $\kappa$ and $\theta$ as well and not keep them fixed. Tests with alternative values of $\kappa$ and $\theta$ resulted in different volatility parameters and different fits, suggesting that calibrating these parameters might improve pricing accuracy.

# A Appendix

## A.1 Details for the Solutions of the Riccati Equations and Piecewise Constant Parameters

In equation (41), a factor $-(z+1)$ appears in front of the integral inside the expectation. The equation is solved for a general RHS which is denoted $V$ and general boundary conditions $A(T,T) = a$ and $B(T,T) = b$, so that all the cases can be taken care of at once. As a result, the equation for $B$ in (21) is modified to

$$B'_t(t,u) - \kappa B(t,u) - \frac{1}{2}\sigma^2 B^2(t,u) = -V \tag{50}$$

and the corresponding solutions for both $A$ and $B$ are

$$\begin{gathered} B = \frac{(b\kappa - hb + 2V)e^{h(T-t)} - b\kappa - hb - 2V}{(b\sigma^2 - h - \kappa)e^{h(T-t)} - b\sigma^2 + \kappa - h} \\ A = a + \frac{\kappa\theta}{\sigma^2}((\kappa - h)(T-t) - 2\log((((\kappa - b\sigma^2)(1 - e^{-h(T-t)}) + h(e^{-h(T-t)} + 1)))/(2h))) \\ h = \sqrt{2V\sigma^2 + \kappa^2} \end{gathered} \tag{51}$$

## A.2 Simulation

Numerical simulations were conducted to verify that the implementation and the pricing routine behaved as intended. This was done by fixing the model parameters and then simulating interest rate paths according to the CIR++ dynamics. The probability density of future values conditional on the current value is known for the CIR process and follows a non-central chi-squared distribution (Cox et al. (1985)). Since the distribution is known, it is possible to perform exact simulation. Based on these simulated paths, caplet prices were computed using expression (28). The cap prices was then obtained as the sum of the prices of its constituent caplets. Finally, estimates of the parameters were obtained through calibration and compared with the parameter values used in the simulations.

## A.3 Piecewise Constant Volatility Specification

This section specifies how the piecewise constant volatility parameters are defined over the full maturity span of the caps. The calibrated volatility parameters from the results are also reported here.

Table 6 shows how the volatility parameters are defined over the intervals corresponding to the tenor structure of the caps.

| Interval | Volatility Parameter |
|---|---|
| $[0, 1]$ | $\sigma_1$ |
| $[1, 2]$ | $\sigma_2$ |
| $[2, 3]$ | $\sigma_3$ |
| $[3, 4]$ | $\sigma_4$ |
| $[4, 5]$ | $\sigma_5$ |
| $[5, 6]$ | $\sigma_6$ |
| $[6, 7]$ | $\sigma_7$ |
| $[7, 8]$ | $\sigma_8$ |
| $[8, 9]$ | $\sigma_9$ |
| $[9, 10]$ | $\sigma_{10}$ |
| $[10, 12]$ | $\sigma_{11}$ |
| $[12, 15]$ | $\sigma_{12}$ |
| $[15, 20]$ | $\sigma_{13}$ |
| $[20, 30]$ | $\sigma_{14}$ |

Table 6: Piecewise constant volatility parameters and corresponding maturity intervals given in years.

Figure 15 shows the calibrated volatility parameters for the different strike sets used in section 4.3 for Periods 1 and 2.

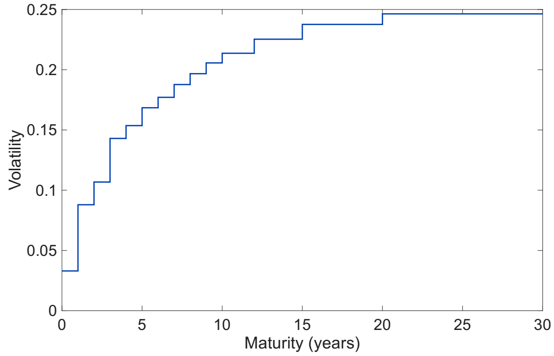


(a) Strike set 1, period 1.

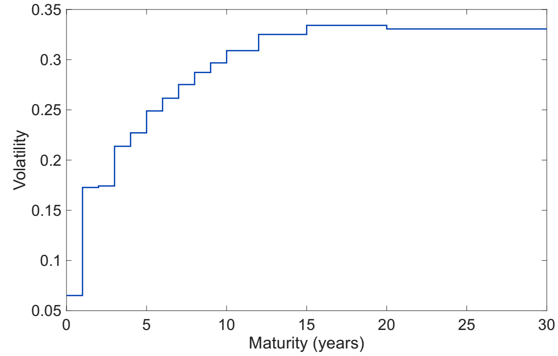


(b) Strike set 1, period 2.

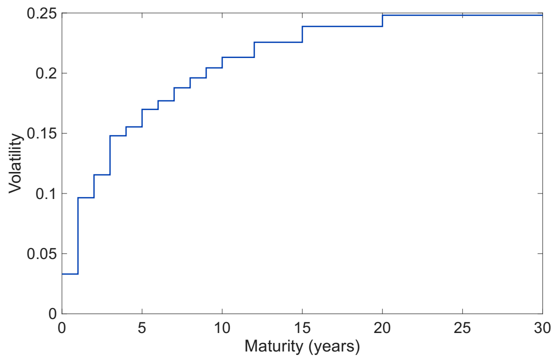


(c) Strike set 2, period 1.

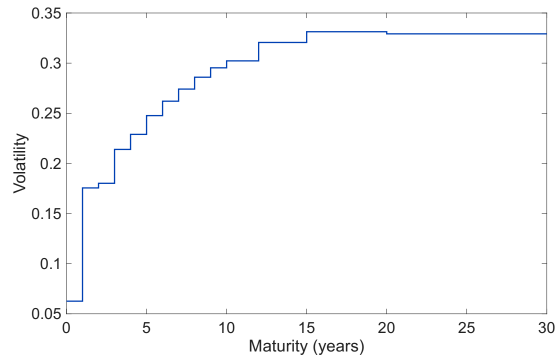


(d) Strike set 2, period 2.

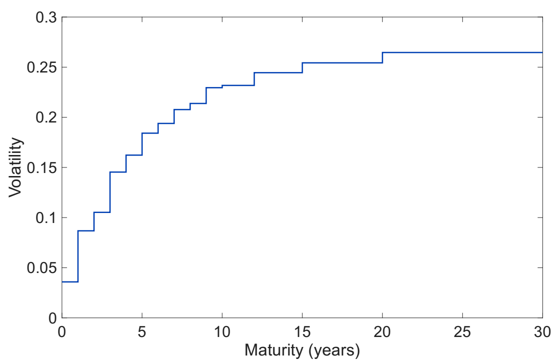


(e) Strike set 3, period 1.

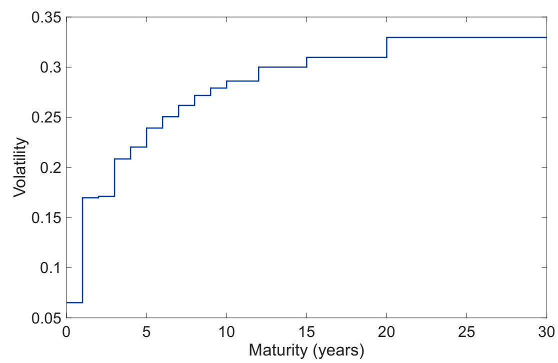


(f) Strike set 3, period 2.

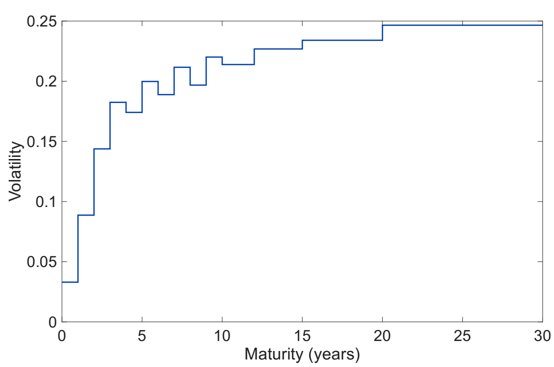


(g) Strike set 4, period 1.

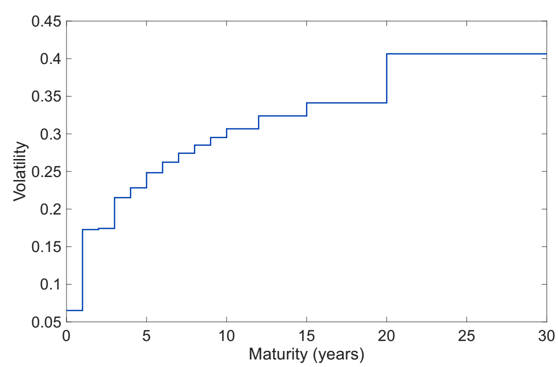


(h) Strike set 4, period 2.

Figure 15: Volatility parameters for the different calibration sets in Period 1 and 2. The calibration sets are listed in Tables 2 and 4.

Figure 16 shows the volatility parameters when calibrated to one day and five days of data before the representative dates in Periods 1 and 2. The parameters are used when modeling the prices in Figures 11a and 12a.

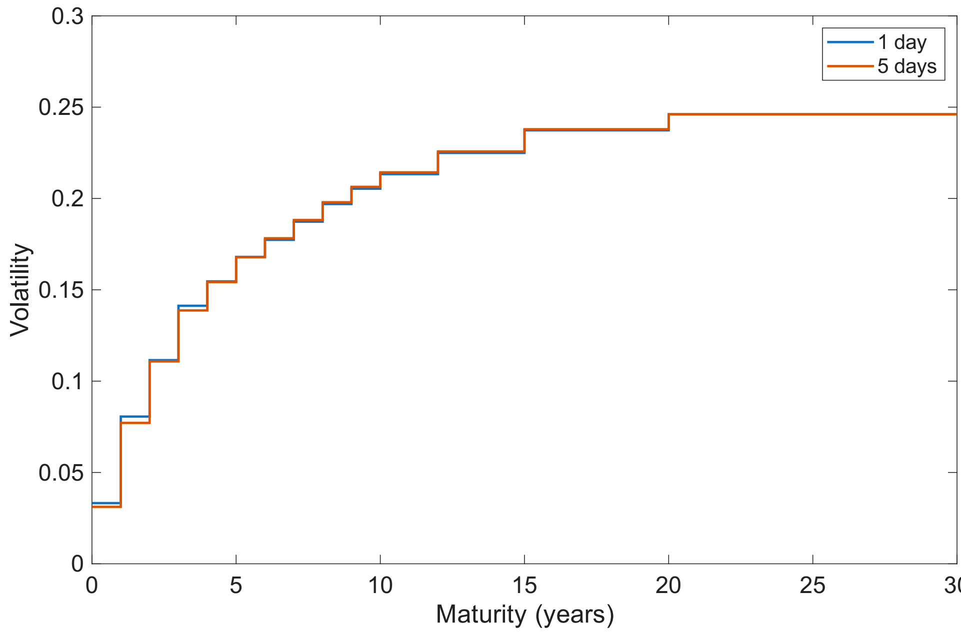


(a) Volatility as a function of maturity for the two calibration sets in Period 1.

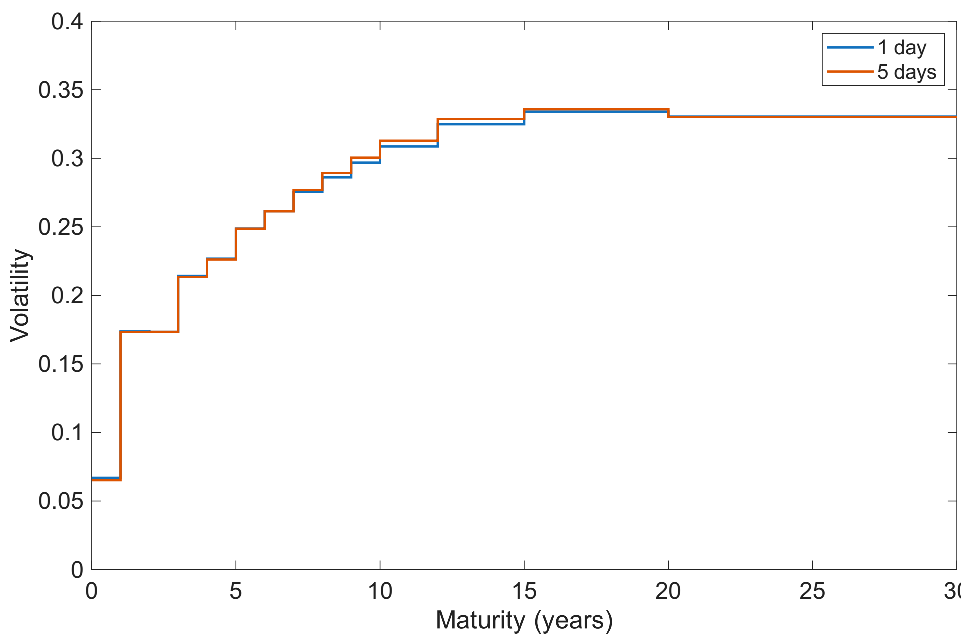


(b) Volatility as a function of maturity for the two calibration sets in Period 2.

Figure 16: Piecewise constant volatility as a function of maturity. The volatility parameters have been calibrated to one day and five days of cap prices before the representative dates in Periods 1 and 2.

## A.4 Strike Specification

This section reports the strikes used in the analysis.

| Date | Maturity | Strike |
|---|---|---|
| Representative Date 1 | 1 | 5.7 |
| | 2 | 18.9 |
| | 3 | 39.0 |
| | 4 | 56.9 |
| | 5 | 70.6 |
| | 6 | 81.8 |
| | 7 | 90.8 |
| | 8 | 98.0 |
| | 9 | 103.8 |
| | 10 | 108.8 |
| | 12 | 116.9 |
| | 15 | 124.8 |
| | 20 | 132.0 |
| | 30 | 133.3 |
| Representative Date 2 | 1 | 507.0 |
| | 2 | 467.8 |
| | 3 | 442.7 |
| | 4 | 426.9 |
| | 5 | 417.2 |
| | 6 | 411.4 |
| | 7 | 407.6 |
| | 8 | 405.1 |
| | 9 | 403.6 |
| | 10 | 402.6 |
| | 12 | 401.9 |
| | 15 | 401.5 |
| | 20 | 397.6 |
| | 30 | 378.7 |

Table 7: ATM strikes per maturity for the two representative dates in Table 1. The representative dates 1 and 2 correspond to the representative dates in Period 1 and 2, respectively. The strikes are given in basis points and the maturities are given in years.

# References


Alvarez, L. H. (1998). Zero coupon bonds and affine term structures: reconsidering the one-factor model. *Insurance: Mathematics and Economics*, 23(1):85–90.

Bank for International Settlements (2025a). Otc derivatives statistics at end-june 2025.

Bank for International Settlements (2025b). Otc interest rate derivatives turnover in april 2025. Triennial Central Bank Survey.

Björk, T. (2009). *Arbitrage Theory in Continuous Time*. Oxford university press.

Brigo, D. and Mercurio, F. (2006). *Interest Rate Models - Theory and Practice*. Springer, 2nd edition.

Choi, J., Kwak, M., Tee, C. W., and Wang, Y. (2022). A Black–Scholes user's guide to the Bachelier model. *Journal of Futures Markets*, 42(5):959–980.

Cox, J. C., Ingersoll, J. E., Ross, S. A., et al. (1985). A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407.

Federal Reserve Bank of New York (n.d.). Secured overnight financing rate data.

Huan, X., Previts, G. J., and Parbonetti, A. (2023). Understanding the libor scandal: the historical, the ethical, and the technological. *Journal of Banking Regulation*, 24(4):403–419.

Lindström, E., Ströjby, J., Brodén, M., Wiktorsson, M., and Holst, J. (2008). Sequential calibration of options. *Computational Statistics & Data Analysis*, 52(6):2877–2891.

Lyashenko, A. and Mercurio, F. (2019). Looking forward to backward-looking rates: a modeling framework for term rates replacing libor. *Available at SSRN 3330240*.

Stoltz, R. and Sönne, F. (2025). Modeling the post-LIBOR interest rate market-calibration of a time-dependent Hull-White model to SOFR caps. Master's thesis, Centre for Mathematical Sciences, Lund University, Lund, Sweden. Available at https://lup.lub.lu.se/student-papers/search/publication/9201629.