POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

A new Model for the Joint Calibration of SPX and VIX derivatives

Tesi di Laurea Magistrale in
Mathematical Engineering - Ingegneria Matematica

Author: **Matteo Algeri**

Student ID: 225647
Advisor: Prof. Roberto Baviera
Co-advisors: Prof. Michele Azzone
Academic Year: 2024-25

Abstract

The VIX index has become a central benchmark for market-implied volatility, serving both as an indicator of investor sentiment and as a tradable instrument through futures and options. As the VIX derivatives market has grown, so has the need for models that can jointly capture the dynamics of both the equity market—via SPX options—and the volatility market—via VIX futures and options.

To address the persistent challenges of jointly calibrating equity and volatility markets, this thesis investigates three continuous-time stochastic volatility models: the classical Heston model, the Ornstein-Uhlenbeck model, and a second-order Continuous Autoregressive model (CAR(2)), which is introduced here as a novel application to equity volatility modeling. Each model is embedded in a Monte Carlo simulation-free joint calibration framework that first aligns key market quantities—such as discount curves, SPX forwards, and VIX futures—before calibrating to SPX and VIX options.

We derive semi-analytic pricing formulae for VIX derivatives and employ Fourier-based methods (FFT) for SPX option pricing. In the CAR(2) model, we improve computational efficiency by adapting the affine ODE framework of Duffie et al. (2000) to solve the model-specific characteristic function. Additionally, this work extends the theoretical insights of Mechkov (2015) by explicitly deriving the characteristic function of the Heston model in the fast mean-reversion limit and establishing a parameter mapping to the Normal Inverse Gaussian distribution.

Overall, this thesis aims to achieve a stable calibration with a compact set of parameters, offering a practical balance between accuracy, tractability, and parsimony.

Keywords: VIX, CAR, joint calibration, stochastic volatility.

Abstract in lingua italiana

L'indice VIX è diventato un punto di riferimento centrale per la volatilità implicita di mercato, fungendo sia da indicatore della fiducia degli investitori sia da strumento negoziabile tramite futures e opzioni. Con la crescente diffusione dei derivati sul VIX, è emersa l'esigenza di modelli capaci di catturare congiuntamente la dinamica del mercato azionario—attraverso le opzioni sull'indice SPX—e del mercato della volatilità—attraverso futures e opzioni sul VIX. Per affrontare le persistenti difficoltà legate alla calibrazione congiunta di questi due mercati, questa tesi analizza tre modelli di volatilità stocastica in tempo continuo: il modello classico di Heston, il modello di Ornstein–Uhlenbeck e un modello autoregressivo continuo di secondo ordine (CAR(2)), proposto qui come nuova applicazione alla modellizzazione della volatilità azionaria. Ciascun modello è inserito in un contesto di calibrazione congiunta privo di simulazioni Monte Carlo, che prevede un allineamento iniziale di grandezze di mercato chiave—come curve di sconto, forward SPX e futures VIX—seguito dalla calibrazione su opzioni SPX e VIX. Deriviamo formule di prezzo semi-analitiche per i derivati sul VIX e impieghiamo metodi basati sulla trasformata di Fourier (FFT) per la valutazione delle opzioni su SPX. Nel modello CAR(2) miglioriamo l'efficienza computazionale adattando il sistema affine di equazioni differenziali di Duffie et al. (2000) per calcolare la funzione caratteristica del modello. Inoltre, questo lavoro estende i risultati teorici di Mechkov (2015) derivando esplicitamente la funzione caratteristica del modello di Heston nel caso limite di elevata velocità di ritorno alla media, e identificando una mappatura diretta dei parametri con quelli della distribuzione Normale Inversa Gaussiana. Nel complesso, questa tesi si propone di ottenere una calibrazione stabile con un numero ridotto di parametri, offrendo un compromesso efficace tra accuratezza empirica, trattabilità analitica e parsimonia modellistica.

Parole chiave: VIX, CAR, calibrazione congiunta, volatilità stocastica.

Contents

Introduction

The VIX index, often referred to as Wall Street's "fear gauge", plays a central role in modern financial markets. It measures the market's expectation of thirty-day forward-looking volatility of the S&P 500 (SPX) index, derived from a wide range of real-time SPX option prices. Initially introduced in 1993 by the Chicago Board Options Exchange (Cboe) and refined in 2003 to better capture market conditions, the VIX has become the standard benchmark for equity market volatility. A liquid derivatives market has developed around it, including VIX futures and options, which are widely used to hedge volatility risk in option portfolios independently of directional exposure to the SPX itself.

SPX options are written directly on the equity index, while VIX derivatives are based on market's expectation of future volatility. This structural difference poses a significant modeling challenge, and capturing the dynamics of both markets with a consistent stochastic volatility framework is nontrivial.

This thesis addresses this challenge by developing a joint calibration framework for SPX and VIX derivatives. Specifically, we implement two established continuous-time stochastic volatility models and introduce a new model designed to more effectively capture the autoregressive and mean-reverting nature of volatility dynamics. Calibration is performed through a multi-stage cascade: we first align market-implied quantities such as discount factors, forward levels and VIX volatility, and subsequently fit the SPX smile using a minimal set of parameters. This approach improves both the interpretability and the internal consistency of the joint calibration process.

Main Goals

The main goals and key contributions of this thesis can be summarized as follows:

- **Implementation of Heston and OU stochastic volatility models:** Study the Heston and Ornstein-Uhlenbeck models in the context of joint calibration of SPX and VIX derivatives. This includes analyzing how VIX options can be priced within these continuous-time frameworks and extending the theoretical work of Mechkov (2015) by establishing a connection between the fast mean-reversion limit of the Heston model and the Normal Inverse Gaussian process.
- **Introduction of CAR(2) model for joint SPX and VIX calibration:** Introduce and develop a novel application of the second-order Continuous Autoregressive (CAR(2)) process to volatility modeling. Leveraging the continuity of the process and its affine structure, we derive closed-form expressions for VIX derivatives and compute the characteristic function via a system of ODEs. This enables semi-analytic pricing of both SPX and VIX options and supports efficient joint calibration within a simulation-free framework.
- **Model comparison and calibration cascade methodology:** Compare the empirical performance of the three models using real SPX and VIX options market data. Implement a multi-step calibration cascade procedure to ensure consistency across derivatives and study each model's strengths and limitations in reproducing observed market smiles and VIX futures prices.

Outline

This thesis is structured as follows:

1. **Literature Review:** Surveys the existing literature on the joint calibration of SPX and VIX derivatives, identifying key challenges and gaps—particularly in terms of calibration methodology, model parsimony, and tractability. The Chapter also reviews the models used in this thesis—Heston, Ornstein–Uhlenbeck, and CAR(2)—highlighting their theoretical foundations and practical relevance in volatility modeling.
2. **Heston Model:** Presents the classical Heston stochastic volatility model, derives its characteristic function, explores the fast mean-reversion limit connecting it to the Normal Inverse Gaussian (NIG) distribution, and provides formulae for pricing VIX futures and options.

3. **Ornstein-Uhlenbeck Model:** Introduces a stochastic volatility model driven by an OU process. The Chapter derives the characteristic function and provides closed-form expressions for forward variance and VIX derivative pricing, emphasizing the model's tractability and parsimony.

4. **Continuous Autoregressive Model:** Introduces the CAR(2) model and adapts it to the volatility modeling setting. It discusses the derivation of the characteristic function via a system of ODEs and provides semi-analytical pricing formulae for both SPX and VIX derivatives.

5. **Calibration:** Describes the joint calibration methodology based on a multi-stage cascade approach. Market quantities such as discount factors, forwards, and VIX futures are first aligned before fitting SPX implied volatility smiles.

6. **Results:** Evaluates the calibration performance of the Heston, OU, and CAR(2) models across different maturities. Strengths and weaknesses of each model are discussed in terms of pricing accuracy, parameter stability, and computational efficiency.

7. **Conclusions:** Summarizes the main findings, emphasizes the advantages and limitations of the proposed framework, and discusses possible future research directions.

1 | Literature Review

1.1. The Joint Calibration challenge

Since the introduction of VIX options in 2006, many analysts have attempted to build models capable of consistently pricing both SPX options and VIX derivatives—futures and options—within a unified framework.
However, achieving a robust joint calibration of these two markets remains a challenging task.

Early attempts, such as the influential work of Bergomi (2008), introduced the concept of a forward variance curve as an input to model the joint dynamics of SPX and VIX. While effective in some respects, this approach relied heavily on externally supplied volatility curves and Monte Carlo simulations, which limit interpretability and increase computational costs.

Gatheral (2008) provided an important critical assessment of the limitations of continuous-time models in jointly capturing SPX and VIX derivatives, underscoring the structural incompatibility that arises when models fail to produce consistent state-price densities across the two markets. This theoretical tension was further formalized in Song and Xiu (2012), who rigorously showed that no existing continuous diffusion model could simultaneously calibrate both markets with sufficient precision.

To address this issue, more flexible (yet often more complex) models have emerged. The Quintic model of Jaber et al. (2023) introduced high-degree polynomial volatility dynamics, requiring both a forward input curve and Monte Carlo methods. Guyon (2020) developed a comprehensive (but highly intricate) framework involving jumps, local volatility, and a hybrid calibration technique, at the cost of tractability and interpretability.

A relevant academic contribution to the joint calibration literature is the work of Maffei (2020), which presents and compares two distinct approaches to jointly calibrate SPX and VIX derivatives: the Lévy jump-diffusion framework of Cont and Kokholm (2013), and the discrete joint distribution model proposed by Guyon (2020). The work provides

both theoretical insights and practical implementation details, including the use of Fourier methods and Monte Carlo simulations. This empirical analysis on real market data supports the validity of these models and investigates the presence of arbitrage across the SPX and VIX markets, offering valuable evidence of the challenges and trade-offs in balancing model realism and tractability.

Kokholm and Stisen (2015) proposed a joint framework based on a Heston model with jumps, showing that adding jumps to the price process improves the model's ability to capture the market-implied dynamics of both VIX and SPX derivatives. While their model achieved better empirical calibration than purely diffusive counterparts, it still required a significantly more complex calibration procedure and lacked closed-form tractability for certain quantities.

More recently, Cuchiero et al. (2024) proposed a stochastic volatility model with volatility dynamics described by a linear function of the signature of a polynomial diffusion process. Exploiting this structure, they derive closed form expressions for the VIX squared and use Fourier-based pricing for both SPX and VIX options. The model achieves highly accurate joint calibration results without requiring jumps or rough volatility. However, it relies on up to 85 parameters, time-varying inputs, and offline precomputation—resulting in greater complexity and reduced interpretability compared to more parsimonious alternatives.

In contrast to these approaches, this thesis proposes a parsimonious, interpretable, and computationally efficient framework. We implement two established models—the Heston and the Ornstein-Uhlenbeck models—and introduce the CAR(2) model, a continuous-time affine stochastic volatility process with two autoregressive components, into the joint calibration problem. These models provide explicit expressions for VIX futures and options, and use Fourier methods (FFT) for SPX option pricing, entirely avoiding Monte Carlo simulations.

1.2. Heston

The Heston model, first introduced by Heston (1993), is one of the most widely adopted continuous-time stochastic volatility models in financial mathematics. It features a mean-reverting square-root process for variance, allowing for closed-form solutions for European option prices via characteristic function. This tractability, combined with its ability to generate volatility smiles, has made it a standard benchmark in option pricing.

An extension by Heston and Nandi (1998) introduced the continuous-time limit of a discrete-time GARCH formulation and examined its behavior under extreme negative

correlation between asset returns and volatility—a regime that also emerges in this thesis when jointly calibrating SPX and VIX derivatives.

More recently, a theoretical work by Mechkov (2015) explored the limit behavior of the Heston model under infinitely fast mean reversion, showing its convergence to a Normal Inverse Gaussian process. This established a bridge between stochastic volatility and Lévy process-based modeling.

1.3. Ornstein-Uhlenbeck

Originally introduced for modeling physical phenomena such as the velocity of a Brownian particle under friction (see e.g. Uhlenbeck and Ornstein (1930)), the Ornstein–Uhlenbeck (OU) process has found wide use in financial modeling, thanks to its tractable Gaussian dynamics and mean-reverting behavior. One of its most influential applications in finance is found in the Vasicek (1977) model, where it is employed to model the evolution of interest rates.

In the context of stochastic volatility modeling, Schöbel and Zhu (1999) adopted an OU process for the dynamics of instantaneous volatility. By leveraging the Gaussian structure of the process, they derived semi-analytical pricing formulas for European options using characteristic function techniques. This framework offers a tractable alternative to more intricate stochastic volatility models—such as those involving jumps or multiple volatility factors—while still capturing key features of implied volatility surfaces observed in equity markets.

1.4. Continuous Autoregressive

Although Continuous-time Autoregressive (CAR) and CARMA processes have been employed in various areas of quantitative finance—such as the modeling of temperature derivatives and electricity spot prices, as in Benth et al. (2008) and the thesis of Engan (2013)—their application to equity volatility modeling has received little attention. An exception is found in the context of non-Gaussian pure jump Lévy-driven CARMA models for stochastic volatility, where Tauchen and Todorov (2004) developed efficient simulation schemes that highlight their flexibility in capturing persistent volatility dynamics. However, these volatility-related works mostly focus on simulation-based estimation rather than analytical pricing or calibration.

The theoretical foundation for Continuous-time Autoregressive Moving Average (CARMA) processes was laid out in Brockwell (2001a) and Brockwell (2001b), which introduced

Lévy-driven CARMA models as a flexible class of stochastic processes for modeling complex temporal dynamics in continuous time, and discussed their properties and connections to discrete-time ARMA processes.

This thesis, by contrast, extends the CAR framework in a novel direction—applying it to the joint calibration of SPX and VIX derivatives through the introduction of a second-order CAR (CAR(2)) process for equity volatility modeling in continuous time. Leveraging the model's Gaussian and affine structure, we derive semi-analytical pricing formulas for SPX and VIX options, and implement a coherent calibration procedure based on Fourier pricing techniques. By successfully implementing and applying the CAR(2), this thesis opens the door to further exploration of continuous-autoregressive volatility structures in the joint calibration framework.

2 | Heston Model

2.1. Introduction

The Heston model, introduced in Heston (1993), is one of the most influential stochastic volatility models in financial mathematics. It effectively captures important market features such as the implied volatility smile, making it widely adopted both in academic research and industry practice. In this Chapter, we present the Heston model as a first benchmark within our joint calibration framework.

We begin by introducing the model dynamics and deriving the characteristic function of the log-price process. This result is fundamental for semi-analytical option pricing using Fourier-based methods. We then explore the limit case where the volatility mean-reversion speed tends to infinity, following the work of Mechkov (2015). This regime allows us to connect the Heston model with the Normal Inverse Gaussian (NIG) process and establish a direct and practical mapping between their parameters.

Lastly, we derive the pricing formulae for VIX derivatives—futures and options—under the Heston model. Although semi-analytical expressions are available, they involve the noncentral Chi-squared distribution, which increases computational complexity. These derivations establish the foundation for assessing the model's calibration performance in later Sections.

2.2. Model

Mechkov (2015) introduces an alternative formulation of the classical Heston model, which rewrites the variance process in a normalized and rescaled form. We write the standard Heston model that follows the Cox–Ingersoll–Ross (CIR) dynamics (Cox et al. (1985)) as:

$$\begin{cases} dx_t & = -\frac{1}{2} v_t dt + \sqrt{v_t} dW_x \ , \\ dv_t & = a\left(\theta - v_t\right) \, dt + \chi \sqrt{v_t} \, dW_v \ , \\ < dW_x dW_v > & = \rho dt \ , \end{cases} \tag{2.1}$$

where $x_t = \log\left(\frac{F(t,T)}{F(t_0,T)}\right)$, v_t is the instantaneous variance, a is the speed of mean reversion, θ the long-term mean, χ the volatility of volatility, and ρ the correlation between the Brownian motions. The risk-free rate r is not explicit, as we directly model the forward price process $F(t,T)$ to ensure the martingale property under the risk-neutral measure.

Mechkov reformulates this model by rescaling the variance process and introducing a volatility factor z_t with unitary mean. The resulting dynamics are:

$$\begin{cases} dx_t & = -\frac{1}{2} z_t \sigma^2 \, dt + \sigma \sqrt{z_t} \, dW_x \quad , \\ dz_t & = a\left((1 - z_t)\, dt + \gamma \sqrt{z_t} \, dW_z\right) \quad , \\ < dW_x, dW_z > & = \rho dt \quad , \end{cases} \tag{2.2}$$

where z_t is a mean-reverting process that scales the constant volatility parameter σ. Here, γ plays the role of relative volatility of the factor z_t, and the process reverts to a fixed unitary mean.

The two formulations are equivalent under suitable parameter transformations. The second, normalized formulation is especially useful for analytical and asymptotic analysis, such as in the infinite-reversion speed limit discussed later in this Chapter.

2.3. Characteristic Function

To compute the characteristic function of the log-return distribution $\phi_{x_t}(u)$ of the process defined in Equation (2.2), we consider the following conditional expectation:

$$f_t(x,z) = \mathbb{E}[e^{ipx_T} | x_t = x, z_t = z] \quad , \tag{2.3}$$

where $t \leq T$, and we set $x_0 = 0$.
Applying Ito lemma to $f_t(x,z)$ and imposing the martingale condition, leads to the following backward Kolmogorov PDE:

$$\frac{\partial f}{\partial t} = \frac{1}{2}\sigma^2 z \frac{\partial f}{\partial x} - a(1-z)\frac{\partial f}{\partial z} - \frac{1}{2}\sigma^2 z \frac{\partial^2 f}{\partial^2 x} - \frac{1}{2}a^2\gamma^2 z \frac{\partial^2 f}{\partial^2 z} - a\rho\gamma\sigma z \frac{\partial^2 f}{\partial x \partial z} \quad , \tag{2.4}$$

with terminal condition at $t = T$ given by:

$$f_T(x,z) = e^{ipx} \quad .$$

We assume a solution of the exponential-affine form:

$$\begin{cases} f_t(x,z) &= e^{A(t)+B(t)x+a^{-1}C(t)z} \ , \\ B(t) &= ip \ , \end{cases} \tag{2.5}$$

where $A(t)$, $B(t)$ and $C(t)$ are deterministic functions of time t, and a is the speed of the mean reversion of the process.
We compute the partial derivatives required in Equation (2.4), simplifying the notation by setting $A(t) = A$, $B(t) = B$, $C(t) = C$:

$$\begin{cases} \dfrac{\partial f}{\partial t} &= f[A' + B'x + a^{-1}C'z] \ , \\ \dfrac{\partial f}{\partial x} &= fB \ , \\ \dfrac{\partial f}{\partial z} &= fa^{-1}C \ , \\ \dfrac{\partial^2 f}{\partial x^2} &= fB^2 \ , \\ \dfrac{\partial^2 f}{\partial z^2} &= fa^{-2}C^2 \ , \\ \dfrac{\partial^2 f}{\partial x \partial z} &= fBa^{-1}C \ . \end{cases}$$

So, the PDE defined in Equation (2.4) reduces to the following system of ODEs:

$$\begin{cases} \dfrac{dA}{dt} = -C \ , \\ \dfrac{dB}{dt} = 0 \ , \\ \dfrac{dC}{dt} = C^2(-\dfrac{1}{2}\gamma^2 a) + C(-Ba\gamma\sigma\rho + a) + \dfrac{1}{2}a\sigma^2 B(1-B) \ , \end{cases} \tag{2.6}$$

subject to the terminal conditions:

$$A(T) = 0 \ , \qquad B(T) = iu \ , \qquad C(T) = 0 \ .$$

The ODE for $C(t)$ is a Riccati equation, and its solution is given by:

$$C(t) = \frac{-\gamma^{-1}\Omega(1-\Psi^2)}{\coth\left\{\frac{1}{2}a\gamma\Omega(T-t)\right\} + \Psi} \ , \tag{2.7}$$

where the functions Ω and Ψ are defined as follows:

$$\begin{cases} \Omega = \sqrt{\sigma^2 B(1-B) + (\gamma^{-1} - \rho\sigma B)^2} \\ \Psi = \dfrac{\gamma^{-1} - \rho\sigma B}{\Omega} \end{cases} .$$

A complete derivation of this expression is provided in Appendix A.1.

Solving the system in (2.6) yields the Heston characteristic function, expressed as:

$$\phi_T(u) = \frac{\exp\left\{\frac{-iu\sigma\rho\gamma+1}{\gamma^2}T\right\} \exp\left\{\frac{-\gamma^{-1}\sigma^2 iu(1-iu)\Omega^{-1}}{\coth\left(\frac{a\gamma\sigma}{2}T\right)+\Psi}(a^{-1}z_0)\right\}}{\left[\cosh\left(\frac{a\gamma\sigma}{2}T\right) + \frac{1-iu\sigma\rho\gamma}{\gamma\Omega}\sinh\left(\frac{a\gamma\sigma}{2}T\right)\right]^{\frac{2}{a\gamma^2}}} . \tag{2.8}$$

A detailed derivation is provided in Appendix A.2.

2.3.1. Infinite reversion speed

As shown in Mechkov (2015), we can analyze the Heston model in the limit of infinitely fast mean reversion. Specifically, we study the behavior of the solution (2.7) as $a \to +\infty$. Noting that:

$$\lim_{a\to\infty} \coth\left(\frac{1}{2}a\gamma\Omega(T-t)\right) = 1 ,$$

the coefficient $C(t)$ converges to its stationary value $\overline{C}_{\sigma,\gamma,\rho}$:

$$\lim_{a\to\infty} C(t) = \overline{C}_{\sigma,\gamma,\rho}(u) = \gamma^{-2}\left(1 - iu\sigma\gamma\rho - \sqrt{(\sigma\gamma)^2(u^2+iu) + (1-iu\sigma\gamma\rho)^2}\right) , \tag{2.9}$$

which is independent of t.
Substituting the limit (2.9) into the solution of the backward PDE (2.4), the function $f_t(x,z)$ simplifies to:

$$f_t(x,z) = \exp\left\{(T-t)\overline{C}_{\sigma,\gamma,\rho}(p) + ipx + \overline{C}_{\sigma,\gamma,\rho}(p)a^{-1}z\right\} . \tag{2.10}$$

In this setting, the stationary distribution of the variance factor Z follows a Gamma distribution $Z \sim \Gamma\left(\frac{2}{a\gamma^2}, \frac{a\gamma^2}{2}\right)$, so that $a^{-1}Z \sim \Gamma\left(\frac{2}{a\gamma^2}, \frac{\gamma^2}{2}\right)$, where the parameter a of the speed of the mean-reversion process just affects the shape parameter of the Gamma distribution.

This scaled random variable has the following:

$$\begin{cases} \mathbb{E}[a^{-1}Z] = \dfrac{1}{a} \ , \\ Var[a^{-1}Z] = \dfrac{\gamma^2}{2a} \ , \end{cases} \qquad as \quad a \to +\infty \ . \tag{2.11}$$

Since the variance tends to zero, $a^{-1}Z$ collapses to its mean, i.e. zero.
By evaluating the function $f_t(x, z)$ at $x_o = 0$ and $z = z_0$, the limit characteristic function becomes:

$$\phi_T(u) = f_0(0, z_0) = \exp\Big\{T\overline{C}_{\sigma,\gamma,\rho}(u)\Big\} \ . \tag{2.12}$$

Two conclusions naturally follow.
First, the process x_t becomes Markovian with no dependence on the factor z_t, indicating that volatility fluctuations are absorbed into the marginal distribution.
Second, the characteristic function is exponential in T, implying that x_t becomes a Lévy process with independent and stationary increments. This convergence establishes a formal link between the fast-reverting Heston model and an infinitely divisible process, specifically the Normal Inverse Gaussian (NIG), as discussed in Mechkov (2015), and further studied in the next Section.

2.3.2. Link to Normal Inverse Gaussian

As demonstrated in Mechkov (2015) and showed in the previous Section, the characteristic function of the Heston model in the fast mean-reversion limit (FRH) coincides with that of a Normal Inverse Gaussian (NIG) process. In this Section, we establish explicit mappings between the parameters of the FRH model and two popular formulations of the NIG process:

- **FRH**: Fast-Reversion Heston's model, with characteristic function in (2.12);
- **NIG$_3$**: NIG process in its affine representation that can be found in Equation 15 of Azzone and Baviera (2023) as

$$\phi_t(u; \tilde{\sigma}, \tilde{\eta}, \tilde{k}) = \mathcal{L}\left(iu\left(\frac{1}{2} + \tilde{\eta}\right)\tilde{\sigma}^2 + \frac{u^2\tilde{\sigma}^2}{2}; \tilde{k}\right) e^{-iu \log \mathcal{L}_t(\tilde{\eta}\tilde{\sigma}^2; \tilde{k})} \ ,$$

  with $\tilde{\sigma}$, $\tilde{k}$ continuous on $[0, \infty)$ and $\tilde{\eta}$ continuous on $(0, \infty)$, with $\tilde{\sigma} > 0$, $\tilde{k}, \tilde{\eta} \geq 0$, and Laplace exponent

$$\ln \mathcal{L}_t(w; \tilde{k}) := \frac{t}{\tilde{k}}\left\{1 - \sqrt{1 + 2w\,\tilde{k}}\right\} \ ;$$

- **NIG**$_4$, NIG in its classical four-parameter parametrization, as in Equation 2.6 of Barndorff-Nielsen (1997):

$$\phi_t(u;\alpha,\beta,\mu,\delta) = \exp\left[i\mu u + \delta\left(\sqrt{\alpha^2-\beta^2} - \sqrt{\alpha^2-(\beta+iu)^2}\right)\right] \quad ,$$

  with $0 \leq |\beta| \leq \alpha, \mu \in \mathbb{R}$ and $0 < \delta$.

We now present Tables 2.1 and 2.2 summarizing the parameter transformations between the models.

Mapping between NIG$_4$ and FRH

NIG$_4(\mu,\ \delta,\ \alpha,\ \beta)$	**FRH**$(\gamma,\ \rho,\ \sigma)$
$\overline{\mu} = \dfrac{\mu}{T} = -\dfrac{\sigma\rho}{\gamma}$	$\sigma^2 = -2\left(\overline{\mu} + \overline{\delta}\dfrac{\beta}{\sqrt{\alpha^2-\beta^2}}\right)$
$\overline{\delta} = \dfrac{\delta}{T} = \dfrac{\sigma\sqrt{1-\rho^2}}{\gamma}$	$\rho = -\overline{\mu}\dfrac{\gamma}{\sigma}$
$\beta = -\dfrac{1}{2}\dfrac{\gamma\sigma - 2\rho}{\sigma\gamma(1-\rho^2)}$	$\gamma^2 = \dfrac{1}{\overline{\delta}\sqrt{\alpha^2-\beta^2}}$
$\alpha = \dfrac{1}{2}\dfrac{\sqrt{4-4\rho\gamma\sigma+4}}{\gamma\sigma(1-\rho^2)}$	

Table 2.1: Parameters mapping between NIG$_4$ and FRH.

Mapping between FRH and NIG$_3$

FRH$(\gamma,\ \rho,\ \sigma)$	**NIG**$_3(\tilde{\sigma},\ \tilde{\eta},\ \tilde{k})$
$\rho^2 = -\dfrac{\tilde{k}\tilde{\sigma}^2}{\tilde{k}\tilde{\sigma}^2(1+2\tilde{\eta}) + 2 - 2\sqrt{1+2\tilde{k}\tilde{\eta}\tilde{\sigma}^2}} + 1$	$\tilde{\eta} = \dfrac{\sigma\gamma - 2\rho}{2\sigma\gamma(1-\rho^2)} - \dfrac{1}{2}$
$\sigma^2 = \dfrac{2 + 2\tilde{k}\tilde{\eta}\tilde{\sigma}^2 - 2\sqrt{1+2\tilde{k}\tilde{\eta}\tilde{\sigma}^2}}{\tilde{k}\rho^2}$	$\tilde{\sigma}^2 = \sigma^2(1-\rho^2)$
$\gamma^2 = \tilde{k}$	$\tilde{k} = \gamma^2$

Table 2.2: Parameters mapping between FRH and NIG$_3$.

This connection complements and extends the asymptotic analysis of Mechkov (2015),

providing an explicit parametric relationships useful for model calibration and interpretation.

2.4. Pricing VIX derivatives

In this Section, we derive VIX futures and call options pricing formulae under the Heston model, following the approach outlined in Jaber et al. (2023).

As derived by Cuchiero et al. (2024) in Equation 5.1, in continuous time, the squared VIX index at time T can be expressed as:

$$VIX_T^2 \;=\; -\frac{2}{\Delta}\mathbb{E}[\log(S_{T+\Delta}/S_T)|\mathcal{F}_T]\cdot 100^2 \;=\; \frac{100^2}{\Delta}\int_T^{T+\Delta}\xi_T(u)du \;\;, \tag{2.13}$$

where $\Delta = 30$ days (expressed as a year fraction), S_T is the SPX price process at time $T > 0$, and $\xi_T(u)$ denotes the Forward Variance curve at time T for maturity u, with $T \leq u$. This process is defined as:

$$\xi_T(u) := \mathbb{E}[\sigma_u^2|\mathcal{F}_T] \;\;. \tag{2.14}$$

Under the Heston dynamics in Equation (2.2), the instantaneous variance is proportional to the stochastic process z_u. Hence:

$$\begin{aligned}\xi_T(u) \;&=\; \mathbb{E}\left[\sigma_u^2|\mathcal{F}_T\right] \;=\; \mathbb{E}\left[z_u|\mathcal{F}_T\right] \\ &=\; z_T e^{-a(u-T)} + 1 - e^{-a(u-T)} \;\;.\end{aligned} \tag{2.15}$$

Plugging this result into Equation (2.13), we get:

$$VIX_T^2 \;=\; \frac{100^2}{\Delta}\frac{1}{a}(e^{-a\Delta}-1)(1-z_T)+\Delta \;\;. \tag{2.16}$$

The VIX index is thus expressed as a deterministic function of the stochastic factor z_T, which—under the model assumptions—follows a noncentral Chi-squared distribution. As a result, pricing VIX derivatives with payoff function V reduces to integrating over the known distribution of z_T:

$$\mathbb{E}\left[V(\mathrm{VIX}_T)\right] \;=\; \mathbb{E}\left[V\left(\sqrt{\mathrm{VIX}_T^2}\right)\right] \;=\; \mathbb{E}\left[V\left(\sqrt{h(z_T)}\right)\right] \;=\; \int_{\mathbb{R}^+} V\left(\sqrt{h(x)}\right) f_{z_T}(x)dx \;\;,$$

where $h(\cdot)$ denotes the polynomial in Equation (2.16) and $f_{z_T}(x)$ is the probability density

function of z_T.

$$z_u|z_T \sim c(u-T)\chi^2_\nu(\lambda(u-T)) \quad \text{with} \quad u \geq T \ ,$$
$$f_{z_T}(x) = \frac{1}{c(t)} \cdot \texttt{ncx2pdf}\left(\frac{x}{c(t)}, \nu, \lambda(t)\right) \ ,$$

with parameters:

$$\begin{aligned} c(t) &= \frac{a\gamma^2}{4}\left(1-e^{-at}\right) \ , \\ \lambda(t) &= \frac{4e^{-at}}{c(t)} z_T \ , \\ \nu &= \frac{4}{a\gamma^2} \ . \end{aligned}$$

By choosing the appropriate payoff function $V(x)$, we can obtain VIX futures ($V(x) = x$) and VIX call options ($V(x) = (x-K)^+$).

This approach yields a semi-analytical expression for VIX derivative prices, avoiding Monte Carlo simulation. However, due to the involvement of noncentral Chi-squared distributions with non-integer degrees of freedom, the required numerical integration can become slow and unstable. These computational challenges impact the calibration speed and limit the practical effectiveness of the Heston model in capturing VIX market features.

2.5. Summary

In this Chapter, we studied the Heston stochastic volatility model in detail, laying out the theoretical and practical tools required for its use in the joint calibration to SPX and VIX derivatives. We derived the characteristic function of the log-price process, which enables efficient SPX option pricing via Lewis (2001) formula.

We also examined the fast mean-reversion limit of the volatility process, demonstrating the model's convergence to the Normal Inverse Gaussian (NIG) process as described by Mechkov (2015). This theoretical connection clarifies how stochastic volatility behaves in the presence of extreme mean reversion and provides a bridge to more tractable Lévy processes.

Finally, we derived the pricing framework for VIX futures and options within the Heston model. While the model allows for closed-form solutions involving noncentral Chi-squared distributions, the presence of non-integer degrees of freedom introduces considerable computational cost. This trade-off between analytical tractability and numerical efficiency is a recurring theme in the evaluation of stochastic volatility models.

3 | Ornstein-Uhlenbeck Model

3.1. Introduction

This Chapter presents the Ornstein–Uhlenbeck (OU) stochastic volatility model as an alternative continuous-time framework for capturing market dynamics. In contrast to the Heston model, which assumes a square-root diffusion for volatility, the OU model treats the instantaneous volatility as a mean-reverting Gaussian process. Its structural simplicity allows for analytical tractability while still capturing essential features of volatility behavior, such as reversion to a long-term mean and stochastic evolution.

We begin by defining the model and solving the stochastic differential equation for the volatility process. This solution enables us to derive the distributional properties of volatility and its square, which play a central role in pricing derivatives on the volatility index. We then proceed to derive the characteristic function of the log-price process, following an analytical approach, exploiting Feynman–Kac theorem. This result allows for efficient pricing of SPX options using Fourier-based methods.

Finally, we exploit the explicit solution of the OU process to compute closed-form expressions for the forward variance and, consequently, for the squared VIX index. These expressions allow us to derive semi-analytical formulae for pricing VIX futures and options by integrating against the well-known Gaussian density of the volatility process, avoiding the need for computationally intensive Monte Carlo methods.

3.2. Model

We model the stochastic volatility dynamics of the log-price of the underlying asset using an Ornstein–Uhlenbeck (OU) type mean-reverting process. Under this framework, the evolution of x_t and its instantaneous volatility y_t is described by the following stochastic differential equations:

$$\begin{cases} dx_t = -\frac{1}{2}y_t^2\sigma^2 dt + \sigma y_t dW_x(t) \ , \\ dy_t = a(\mu - y_t)dt + \gamma dW_y(t) \ , \end{cases} \tag{3.1}$$

with initial conditions $y_0 = y$, $x_0 = 0$ and the two Brownian motions are correlated $d < W_x, W_y >= \rho dt$, with correlation parameter $\rho \in [-1, 1]$.

In this model, the volatility process y_t exhibits mean reversion towards a long-run average level μ at a speed determined by the parameter $a > 0$, while its vol-of-vol is controlled by the parameter $\gamma > 0$. The parameter σ scales the influence of the volatility process y_t on the x_t dynamics, and the correlation ρ, which is usually negative in equity markets, captures the leverage effect.

3.2.1. Volatility distribution

Consider the volatility process y_t from (3.1), and define $z_t := y_t^2$.
In order to find the distribution of z_t, we first analyze the solution of the OU SDE in (3.1):

$$y_t = y_0 e^{-at} + \mu(1 - e^{-at}) + \gamma \int_0^t e^{a(s-t)}\, dW_y(s) \quad . \tag{3.2}$$

Since the stochastic integral in (3.2) is Gaussian with zero mean, the process y_t is normally distributed with:

$$\begin{aligned} \mathbb{E}[y_t] &= y_0 e^{-at} + \mu(1 - e^{-at}) \quad , \\ Var[y_t] &= \gamma^2 Var\left(\int_0^t e^{a(s-t)}\, dW_y(s)\right) = \gamma^2 \int_0^t e^{2a(s-t)}\, ds \\ &= \gamma^2 e^{-2at} \int_0^t e^{2as}\, ds = \frac{\gamma^2}{2a}\left(1 - e^{-2at}\right) \quad . \end{aligned}$$

We can define

$$\begin{cases} \beta_t = y_0 e^{-at} + \mu(1 - e^{-at}) \quad , \\ \gamma_t^2 = \dfrac{\gamma^2}{2}\left(1 - e^{-2at}\right) \quad , \end{cases} \tag{3.3}$$

so that the distribution of y_t is $y_t \sim \mathcal{N}\left(\beta_t; a^{-1}\gamma_t^2\right)$.

As a consequence, the square volatility $z_t := y_t^2$ follows a noncenteral Chi-squared.
Its mean and variance are given by

$$\begin{aligned} \mathbb{E}[z_t] &= \beta_t^2 + a^{-1}\gamma_t^2 \quad , \\ Var[z_t] &= 2\gamma_t^2 a^{-1}\left(\gamma_t^2 a^{-1} + 2\beta_t^2\right) \quad . \end{aligned}$$

This result is crucial for pricing VIX derivatives, as the squared volatility process enters directly into the computation of conditional expectations and of the characteristic function.

3.3. Characteristic Function

In this Section we derive the Characteristic Function of the OU model, following the methodology outlined in Schöbel and Zhu (1999). The goal is to compute the risk-neutral characteristic function of x_t, defined as:

$$\phi(p) := \mathbb{E}^Q \left[\exp\left(ipx_T\right)\right] \quad . \tag{3.4}$$

Using the dynamics from (3.1), we expand x_T as:

$$\begin{aligned}
\phi(p) &= \mathbb{E}^Q \left[\exp\left\{ipx_T\right\}\right] \\
&= \mathbb{E}^Q \left[\exp\left\{ip\left(x_t - \int_t^T \frac{1}{2}\sigma^2 y^2(u)\,du + \int_t^T \sigma y(u)\,dW_x(u)\right)\right\}\right] \\
&= e^{ipx_t}\,\mathbb{E}^Q \Bigg[\exp\Bigg\{ip\Bigg(-\frac{1}{2}\int_t^T \sigma^2 y^2(u)\,du \\
&\quad + \sigma\left(\rho\int_t^T y(u)\,dW_y(u) + \sqrt{1-\rho^2}\int_t^T y(u)\,dW(u)\right)\Bigg)\Bigg\}\Bigg] \quad ,
\end{aligned}$$

where we use the decomposition of a standard Brownian motion:
$dW_x = \rho dW_y + \sqrt{1-\rho^2}dW$ with $dW dW_y = 0$.
Using the Ito isometry

$$\begin{aligned}
\phi(p) &= e^{ipx_t}\,\mathbb{E}^Q \Bigg[\exp\Bigg\{ip\Bigg(-\frac{1}{2}\int_t^T \sigma^2 y^2(u)\,du + \sigma\rho\int_t^T y(u)\,dW_y(u)\Bigg) \\
&\quad + \frac{1}{2}(1-\rho^2)(ip)^2\sigma^2\int_t^T y^2(u)\,du\Bigg\}\Bigg] \\
&= e^{ipx_t}\,\mathbb{E}^Q \Bigg[\exp\Bigg\{\frac{1}{2}ip\left[ip(1-\rho^2)-1\right]\sigma^2\int_t^T y^2(u)\,du \\
&\quad + ip\sigma\rho\int_t^T y(u)\,dW_y(u)\Bigg\}\Bigg] \quad .
\end{aligned} \tag{3.5}$$

We apply Ito lemma to $f(y) = y^2$ to solve the stochastic integral:

$$y^2(T) - y^2(t) = \int_t^T \left[2ay(u)\left(\mu - y(u)\right) + \gamma^2\right]\,du + \int_t^T 2y(u)\gamma\,dW_y(u) \quad .$$

Plugging it into the expression (3.5), we obtain:

$$\begin{aligned}\phi(p) = e^{ipx_t}\mathbb{E}^Q\Bigg[\exp\Bigg\{&\frac{1}{2}ip\left[ip\sigma^2(1-\rho^2) - 1 + \frac{2\rho}{\sigma a^{-1}\gamma}\right]\int_t^T y^2(u)\,du \\ &- \frac{ip\sigma\rho}{2\gamma}\left[y^2(t) + \gamma^2(T-t) - y^2(T)\right] \\ &- \frac{ip\sigma\rho}{\gamma}a\mu\int_t^T y(u)\,du\Bigg\}\Bigg] \\ \phi(p) = e^{ipx_t}\exp&\left\{\frac{-ip\sigma\rho}{\gamma}\left(y^2(t) + \gamma^2(T-t)\right)\right\}g(y,t,T) \quad .\end{aligned}$$

We thus have to compute the following expectation:

$$g(y,t,T) = \mathbb{E}^Q\left[\exp\left\{\int_t^T(-s_1y^2(u) - s_2y(u))\,du\right\}\exp\left\{s_3y^2(T)\right\}\right] \quad . \tag{3.6}$$

By Feynman-Kac formula, g satisfies the following (backward) partial differential equation:

$$\begin{aligned}\frac{\partial g}{\partial t} + a(\mu - y)\frac{\partial g}{\partial y} + \frac{1}{2}\gamma^2\frac{\partial^2 g}{\partial y^2} - (s_1y^2 + s_2y)g = 0 \quad , \\ g(y,T,T) = e^{s_3y^2} \quad .\end{aligned} \tag{3.7}$$

The solution is of the form

$$\begin{aligned}g(y,t,T) &= \exp\left\{\frac{1}{2}A(t,T)y^2(t) + B(t,T)y(t) + C(t,T) + s_3y^2(t)\right\} \\ &= \exp\left\{\frac{1}{2}D(t,T)y^2(t) + B(t,T)y(t) + C(t,T)\right\} \quad ,\end{aligned} \tag{3.8}$$

with $D(t,T) = A(t,T) + 2s_3$. Substituting the derivatives into the equation (3.7), we obtain a system of Riccati-type ODEs:

$$\begin{cases}D' = -\gamma^2D^2 + 2aD + 2s_1 \quad , \\ B' = B(a - \gamma^2D) - a\mu D + s_2 \quad , \\ C' = -\frac{1}{2}\gamma^2B^2 - a\mu B - B - \frac{1}{2}\gamma^2D \quad ,\end{cases} \tag{3.9}$$

with terminal conditions $D(T,T) = 2s_3$, $B(T,T) = 0$ and $C(T,T) = 0$.
The closed-form solutions are given by:

$$D(t,T) = \frac{1}{\gamma^2}\left\{a - \xi_1 \frac{\sinh\left(\xi_1(T-t)\right) + \xi_2\cosh\left(\xi_1(T-t)\right)}{\cosh\left(\xi_1(T-t)\right) + \xi_2\sinh\left(\xi_1(T-t)\right)}\right\} \quad ,$$

$$B(t,T) = \frac{1}{\gamma^2\xi_1}\left\{\frac{(a\mu\xi_1 - \xi_2\xi_3) + \xi_3\left[\sinh\left(\xi_1(T-t)\right) + \xi_2\cosh\left(\xi_1(T-t)\right)\right]}{\cosh\left(\xi_1(T-t)\right) + \xi_2\sinh\left(\xi_1(T-t)\right)} - a\mu\xi_1\right\} \quad ,$$

$$\begin{aligned} C(t,T) = &-\frac{1}{2}\ln\left[\cosh\left(\xi_1(T-t)\right) + \xi_2\sinh\left(\xi_1(T-t)\right)\right] + \frac{1}{2}a(T-t) \\ &+ \frac{a^2\mu^2\xi_1^2 - \xi_3^2}{2\gamma^2\xi_1^3}\left\{\frac{sinh(\xi_1(T-t))}{\cosh\left(\xi_1(T-t)\right) + \xi_2\sinh\left(\xi_1(T-t)\right)} - \xi_1(T-t)\right\} \\ &+ \frac{\xi_3\left(a\mu\xi_1 - \xi_2\xi_3\right)}{\gamma^2\xi_1^3}\left\{\frac{cosh(\xi_1(T-t)) - 1}{\cosh\left(\xi_1(T-t)\right) + \xi_2\sinh\left(\xi_1(T-t)\right)}\right\} \quad , \end{aligned}$$

with auxiliary parameters

$$\begin{aligned} \xi_1 &= \sqrt{\Psi(p) \; - \; 2\Omega(p)a \; + \; a^2} \quad , \\ \xi_2 &= \frac{1}{\xi_1}\left(a \; - \; \Omega(p)\right) \; = \; \frac{\left(a \; - \; \Omega(p)\right)}{\sqrt{\Psi(p) \; - \; 2\Omega(p)a \; + \; a^2}} \quad , \\ \xi_3 &= \mu\left(a^2 \; - \; \Omega(p)a\right) \quad , \end{aligned} \tag{3.10}$$

and where:

$$\begin{aligned} \Omega(p) &= ip\gamma\sigma\rho \quad , \\ \Psi(p) &= -\gamma^2\left[ip\sigma^2(ip(1-\rho^2) - 1)\right] \quad . \end{aligned} \tag{3.11}$$

This completes the derivation of the characteristic function $\phi(p)$, which is used for pricing SPX options via Lewis (2001) formula.

3.4. Pricing VIX derivatives

In this Section, we exploit the continuity of the OU process, to derive an explicit pricing formula for VIX derivatives.

Exploiting solution in Equation (3.2), and recalling that the Forward Variance process is

$$\xi_T(u) \; = \; \mathbb{E}\left[\sigma_u^2|\mathcal{F}_T\right] \; = \; \mathbb{E}\left[y_u^2|\mathcal{F}_T\right] \quad ,$$

we obtain the following expression for $\xi_T(u)$:

$$\xi_T(u) = y_T^2 e^{-2a(u-T)} + \mu^2 \left(1 - e^{-a(u-T)}\right)^2 + \frac{\gamma^2}{2a}(1 - e^{-2a(u-T)}) + 2\mu y_T e^{-a(u-T)} \left(1 - e^{-a(u-T)}\right) \quad .$$

Plugging this into the definition of the squared VIX index, Equation (2.13), we find:

$$\begin{aligned} VIX_T^2 &= \frac{100^2}{\Delta} \int_T^{T+\Delta} \left\{ y_T^2 e^{-2a(u-T)} + \mu^2 \left(1 - e^{-a(u-T)}\right)^2 \right\} du \\ &+ \frac{100^2}{\Delta} \int_T^{T+\Delta} \left\{ \frac{\gamma^2}{2a}(1 - e^{-2a(u-T)}) + 2\mu y_T e^{-a(u-T)} \left(1 - e^{-a(u-T)}\right) \right\} du \\ &= \frac{100^2}{\Delta} \left\{ y_T^2 \frac{1 - e^{-2a\Delta}}{2a} + \mu^2 \left[\Delta - \frac{2}{a} \left(1 - e^{-a\Delta}\right) + \frac{1}{2a} \left(1 - e^{-2a\Delta}\right) \right] \right\} \\ &+ \frac{100^2}{\Delta} \left\{ \frac{\gamma^2 \Delta}{2a} - \frac{\gamma^2}{4a^2} \left(1 - e^{-2a\Delta}\right) + \frac{\mu y_T}{a} \left(1 - e^{-a\Delta}\right)^2 \right\} \quad . \end{aligned}$$

The expression above shows that VIX_T^2 is a polynomial in y_T.
Since y_T is Gaussian, pricing VIX derivatives with a general payoff function V is straightforward by integrating against the normal density. Denoting $\text{VIX}_T^2(\cdot) = h(\cdot)$, we obtain the following pricing formula:

$$\begin{aligned} \mathbb{E}\left[V(\text{VIX}_T)\right] &= \mathbb{E}\left[V\left(\sqrt{\text{VIX}_T^2}\right)\right] = \mathbb{E}\left[V\left(\sqrt{h(y_T)}\right)\right] \\ &= \frac{1}{\sqrt{2\pi\sigma_{y_T}^2}} \int_{\mathbb{R}} V\left(\sqrt{h(x)}\right) exp\left\{ -\frac{1}{2} \frac{(x - m_{y_T})^2}{\sigma_{y_T}^2} \right\} dx \quad . \end{aligned} \tag{3.12}$$

In particular, to price VIX futures, we set $V(x) = x$, and to price VIX call options, we set $V(x) = (x - K)^+$.

3.5. Summary

In this Chapter, we developed the theoretical foundation and pricing framework for the OU-driven stochastic volatility model. Starting from its easily tractable Gaussian dynamics, we derived the closed-form expression for the forward variance and the squared VIX index, allowing for semi-analytical pricing of VIX derivatives. This provides a computationally efficient and transparent method for evaluating volatility-linked payoffs.

We also derived the characteristic function of the log-asset price under the OU volatility specification, following the work of Schöbel and Zhu (1999), which enables fast and accurate pricing of SPX options using Fourier inversion techniques. The solution relies on a set of coupled Riccati-type ordinary differential equations, solved explicitly.

The model's analytical tractability, combined with its avoidance of the numerical challenges associated with noncentral Chi-squared distributions (as in the Heston model), makes it a valuable alternative for joint calibration tasks. In addition to its low computational complexity, the OU framework is highly parsimonious—requiring only a small number of parameters—yet remains effective in capturing key market features. This balance of simplicity and performance demonstrates its appeal in practical applications.

4 | Continuous Autoregressive Model

4.1. Introduction

This Chapter introduces the CAR(2) (Continuous-time Autoregressive of order 2) model, a continuous-time stochastic volatility model developed for the joint calibration of SPX and VIX derivatives. CAR-type processes, particularly CAR(1) and CAR(2), have been previously employed in fields such as energy finance and meteorology, where they have been successfully applied to model smooth mean-reverting dynamics—particularly in the pricing of temperature derivatives, as discussed in Chapter 10 of Benth et al. (2008).

Inspired by the flexibility and tractability of CAR models, we adapt the CAR(2) framework to volatility modeling, leveraging its Gaussian structure and affine properties. This formulation enables the explicit computation of the forward variance and the squared VIX index, thereby facilitating the semi-analytical pricing of VIX derivatives. Unlike the Heston and OU models, the characteristic function of the CAR(2) process is not available in closed form; rather, it is derived as the numerical solution of a system of ordinary differential equations. While this introduces a degree of computational complexity, the model preserves semi-analytical tractability and supports efficient option pricing through Fourier inversion techniques such as the Lewis (2001) formula. These features position CAR(2) as a promising alternative to traditional stochastic volatility models, offering a balanced trade-off between expressiveness, parsimony, and computational efficiency in the context of joint calibration to SPX and VIX derivatives.

4.2. Model

We consider a continuous-time stochastic volatility model where the log-price dynamics of the underlying asset are driven by a volatility factor y_t that follows a second-order

autoregressive process. Specifically, the process x_t evolves as:

$$dx_t = -\frac{1}{2}y_t^2\sigma^2 dt + \sigma y_t dW_x(t) \ , \tag{4.1}$$

where $\sigma > 0$ is the vol-of-vol parameter, and $W_x(t)$ is a standard Brownian motion. The volatility process y_t is modeled by a linear projection of a process X_t, i.e:

$$\begin{cases} y_t = b^T X_t \ , \\ dX_t = AX_t dt + \gamma e_2 dW(t) \ , \end{cases} \tag{4.2}$$

with $\gamma > 0$ and $W(t)$ a one-dimensional Brownian motion, correlated with $W_x(t)$:

$$d < W_x, W >= \rho dt \ , \qquad \rho \in [-1, 1] \ .$$

The system's matrix and vectors are specified as:

$$A = \begin{bmatrix} 0 & 1 \\ -\alpha_2 & -\alpha_1 \end{bmatrix} \ , \quad b = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \ , \quad e_2 = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \ . \tag{4.3}$$

This formulation defines a second-order continuous-time autoregressive process CAR(2) for y_t, offering enhanced control over volatility dynamics compared to the simpler OU model studied in Chapter 3.
The conditional solution of the volatility process y_t (4.2), given information at time $T \leq u$ is:

$$y_u = b^T e^{A(u-T)} X_T + \gamma \int_T^u b^T e^{A(u-s)} e_2 dW_x(s) \ , \tag{4.4}$$

which facilitates closed-form expressions in the next Sections.

4.2.1. Link to time series

As shown in detail in Chapter 10 of Benth et al. (2008), there is a strong link between continuous-time CAR(p) processes and discrete-time AR(p) models. This correspondence is particularly useful for capturing the autoregressive behavior of the volatility.
Applying an Euler discretization with unitary time step to the SDE system in equation (4.2) yields a discrete-time representation of the latent process $x_p(t)$, $t = 0, 1, \ldots$ of the form:

$$x_p(t+1) - x_p(t) = -\sum_{q=1}^{p} \alpha_{p-q+1}\, x_q(t)\, dt + \sigma(t)\, \epsilon(t) \ , \tag{4.5}$$

where $\epsilon(t) \sim \mathcal{N}(0,1)$ are i.i.d standard Gaussian random variables, and $\Delta t = 1$ in this discretization.
This discrete formulation establishes a direct correspondence between the coefficients of a CAR(p) process and those of its AR(p) counterpart.
In the case of a CAR(2) model, the relation (4.5) becomes:

$$x_1(t+2) = (2-\alpha_1)x_1(t+1) + (1-\alpha_2-\alpha_1)x_1(t) + \gamma\epsilon(t) \ , \tag{4.6}$$

which clearly exhibits an AR(2)-like structure with parameters derived from the continuous-time system.
Importantly, in this thesis we perform a fully implied calibration of the CAR(2) model directly from market data, without relying on any econometric estimation of the AR coefficients from observed time series. This allows us to ensure internal consistency across derivative pricing.

4.2.2. On the stability of the solution

The stability of a continuous-time CAR(2) model is essential for well-defined dynamics and meaningful option pricing. As shown in Brockwell (2001b), a generic CARMA process with constant coefficients is weakly stationary if and only if all the eigenvalues of its matrix A have strictly negative real parts.

For the CAR(2) used in this thesis, the system dynamics are driven by the matrix A defined in (4.3). The eigenvalues of this matrix are obtained from the characteristic equation:

$$\lambda^2 + \alpha_1\lambda + \alpha_2 = 0 \ ,$$

with solutions

$$\lambda_{1,2} = \frac{-\alpha_1 \mp \sqrt{\alpha_1^2 - 4\alpha_2}}{2} \ .$$

To ensure the eigenvalues have strictly negative real parts, the parameters must satisfy:

$$\alpha_1 > 0 \ \text{ and } \ \alpha_2 > 0 \ .$$

Furthermore, based on the results of Brockwell and Lindner (2009), we require that all eigenvalues λ_j of A are real and distinct to guarantee numerical stability and avoid issues in computing the matrix exponential. This leads to the condition:

$$\alpha_1^2 - 4\alpha_2 > 0 \ . \tag{4.7}$$

These conditions ensure that the system does not exhibit oscillatory or explosive behavior, thereby preserving the model's stability and suitability for practical applications.

4.3. Characteristic Function

In this Section, we derive the characteristic function of the CAR(2) model following the affine framework of Duffie et al. (2000). This approach allows us to express the characteristic function in an exponential-affine form, enabling efficient Fourier-based option pricing techniques such as the one introduced by Lewis (2001).
To apply the affine transform methods, we first extend the original state to incorporate all relevant processes. Specifically, we consider the extended state vector X_{full}:

$$X_{\text{full}} = \begin{bmatrix} z \\ x_1 \\ x_2 \\ v = x_1^2 \\ w = x_1 x_2 \\ r = x_2^2 \end{bmatrix} ,$$

whose dynamics are given by the following stochastic differential equations:

$$\begin{cases} dz_t &= -\frac{1}{2}\,\sigma^2\, x_{1,t}^2\, dt \;+\; \sigma\, x_{1,t}\, dW(t) \;, \\ dx_{1,t} &= x_{2,t} dt \;, \\ dx_{2,t} &= (-\alpha_2 x_{1,t} \;-\; \alpha_1 x_{2,t})\, dt \;+\; \gamma dW'(t) \;, \\ dv_t &= 2 w_t dt \;, \\ dw_t &= (r_t \;-\; \alpha_2 v_t \;-\; \alpha_1 w_t)\, dt \;+\; \gamma x_{1,t} dW'(t) \;, \\ dr_t &= \left(-2\alpha_2 w_t \;-\; 2\alpha_1 r_t \;+\; \gamma^2\right) dt \;+\; 2\gamma x_{2,t} dW'(t) \;, \end{cases} \tag{4.8}$$

with initial conditions $z_0 = z$, $x_{1,0} = x_1$, $x_{2,0} = x_2$, and correlation structure $\text{Corr}(\, dW(t),\, dW'(t)\,) = \rho$.
To express the dynamics in the affine framework, we rewrite the Brownian motion as:

$$W(t) = \rho W(t)' + \sqrt{1-\rho^2} W^{\epsilon}(t) \;,$$

with $W^{\epsilon}(t)$ independent of $W'(t)$. the diffusion matrix $\sigma(x)$ in Duffie et al. (2000) (Equation 2.1, page 1349) becomes:

$$\sigma(x) \;=\; \begin{pmatrix} \sqrt{1-\rho^2}\,\sigma\,x_{1,t} & 0 & \rho\,\sigma\,x_{1,t} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \gamma & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \gamma x_{1,t} & 0 & 0 & 0 \\ 0 & 0 & 2\gamma x_{2,t} & 0 & 0 & 0 \end{pmatrix} ,$$

and the corresponding variance-covariance matrix $\sigma(x)\,\sigma(x)^T$ is explicitly computed as:

$$\sigma(x)\,\sigma(x)^T \;=\; \begin{pmatrix} \sigma^2\,v_t & 0 & \rho\sigma\gamma x_{1,t} & 0 & \rho\sigma\gamma v_t & 2\rho\sigma\gamma w_t \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \rho\sigma\gamma x_{1,t} & 0 & \gamma^2 & 0 & \gamma^2 x_{1,t} & 2\gamma^2 x_{2,t} \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \rho\sigma\gamma v_t & 0 & \gamma^2 x_{1,t} & 0 & \gamma^2 v_t & 2\gamma^2 w_t \\ 2\rho\sigma\gamma w_t & 0 & 2\gamma^2 x_{2,t} & 0 & 2\gamma^2 w_t & 4\gamma^2 r_t \end{pmatrix} .$$

Following the affine transform notation from Duffie et al. (2000), we consider the characteristic function:

$$\phi(u; X_t, t, T) = \mathbb{E}\left[e^{u^T X_{\text{full},T}} \mid X_{\text{full},t}\right] = e^{\alpha(t)+\underline{\beta}(t)^T X_{\text{full},t}} \quad , \tag{4.9}$$

where the functions $\underline{\beta}(t)$ and $\alpha(t)$ solve the complex-valued ODEs:

$$\dot{\underline{\beta}}(t) = -K_1^T \underline{\beta}(t) - \frac{1}{2}\underline{\beta}(t)^T H_1 \underline{\beta}(t) \quad , \tag{4.10}$$

$$\dot{\alpha}(t) = -K_0 \underline{\beta}(t) - \frac{1}{2}\underline{\beta}(t)^T H_0 \underline{\beta}(t) \quad , \tag{4.11}$$

with terminal conditions $\underline{\beta}(T) = u$ and $\alpha(T) = 0$.

For the CAR(2), these equations explicitly translate into the following ODE system:

$$\begin{cases} \dot{\alpha}(t) &= \gamma^2\beta_r \ + \ \dfrac{1}{2}\gamma^2\beta_{x_2} \ , \\ \dot{\beta}_z(t) &= 0 \ , \\ \dot{\beta}_{x_1}(t) &= -\alpha_2\beta_{x_2} \ + \ \rho\sigma\gamma\beta_z\beta_{x_2} \ + \ \gamma^2\beta_{x_2}\beta_w \ , \\ \dot{\beta}_{x_2}(t) &= \beta_{x_1} \ - \ \alpha_1\beta_{x_2} \ + \ 2\gamma^2\beta_{x_2}\beta_r \ , \\ \dot{\beta}_v(t) &= -\dfrac{1}{2}\sigma^2\beta_z \ - \ \alpha_2\beta_w \ + \ \dfrac{1}{2}\sigma^2\beta_z^2 \ + \ \rho\sigma\gamma\beta_z\beta_w \ + \ \dfrac{1}{2}\gamma^2\beta_w^2 \ , \\ \dot{\beta}_w(t) &= 2\beta_v \ - \ \alpha_1\beta_w \ - \ 2\alpha_2\beta_r \ + \ 2\rho\sigma\gamma\beta_z\beta_r \ + \ 2\gamma^2\beta_w\beta_r \ , \\ \dot{\beta}_r(t) &= \beta_w \ - \ 2\alpha_1\beta_r \ + \ 2\gamma^2\beta_r^2 \ , \end{cases} \tag{4.12}$$

with terminal conditions $\underline{\beta}(T) = u$ and $\alpha(T) = 0$, as in equations (4.10) and (4.11). The ODEs solutions are found numerically, applying the fourth-order Runge-Kutta method. A simplified version of the code can be found in Appendix C.

In order to apply Lewis (2001) option pricing formula, we need to make sure that

$$\phi(\xi) \text{ analytic in the horizontal strip } \{\xi \in \mathbb{C} : -1 \leq \text{Im}(\xi) \leq 0\} \ .$$

Since the CAR(2) characteristic function relies on numerical ODE integration, to verify its analyticity, we examine its real and imaginary parts over a wide range of complex arguments for each combination of model parameters considered in the calibration. Figure 4.1 presents an illustrative example using the parameters calibrated on the first maturity slice T_1 from the June 20, 2021 dataset. The plots confirm that the function remains smooth and free of singularities or discontinuities along the real axis. In particular, following Theorem 3.1 in Lukacs (1972), which characterizes the analyticity of a characteristic function in terms of its behavior along the imaginary axis, we verify that the function is analytic in a strip around the origin. This empirical validation ensures the numerical stability of the Fourier inversion methods employed in the Lewis formula.

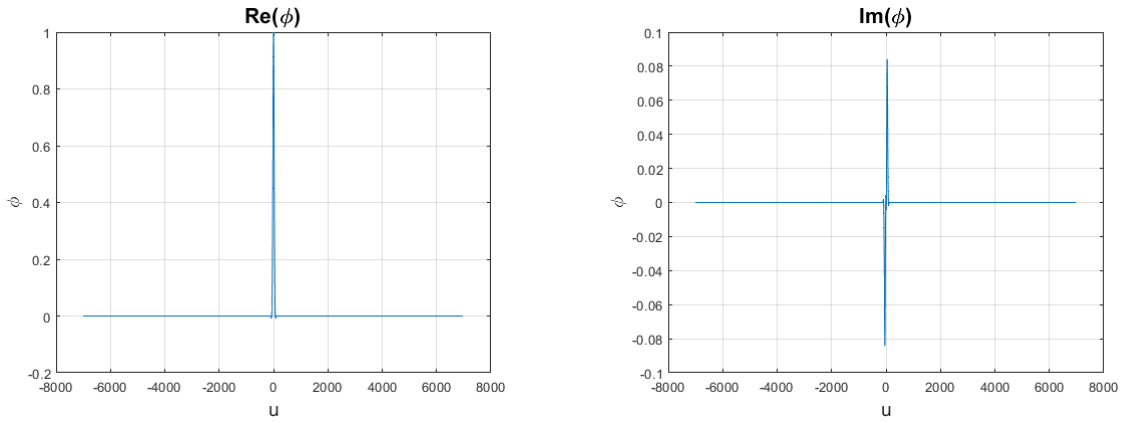

(a) Real part of CAR(2) characteristic function.

(b) Imaginary part of CAR(2) characteristic function.

Figure 4.1: Real and Imaginary parts of CAR(2) characteristic function, with parameter set calibrated on maturity slice T_1 from the June 20, 2021 dataset.

4.4. Pricing VIX derivatives

In this Section, we derive explicit formulae to price VIX derivatives within the CAR(2) model framework, leveraging the continuity of the process.

Let's recall that A is a companion matrix with distinct eigenvalues, hence it is diagonalizable as

$$A \;=\; V\Lambda V^{-1} \;\;, \tag{4.13}$$

where Λ is $\text{diag}\{\lambda_1, \lambda_2\}$ and V is the Vandermonde matrix $(\lambda_1 < \lambda_2)$.

$$V = \begin{pmatrix} 1 & 1 \\ \lambda_1 & \lambda_2 \end{pmatrix} \;\;, \;\; \Lambda = \text{diag}\{\lambda_1, \lambda_2\} = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \;\;, \;\; V^{-1} = \frac{1}{\lambda_2 - \lambda_1} \begin{pmatrix} \lambda_2 & -1 \\ -\lambda_1 & 1 \end{pmatrix} \;\;.$$

Then, as described in Section 4.1 of Engan (2013) and supported by Theorem 1.13 of Higham (2008), the exponential matrix $e^{A\tau} = Ve^{\Lambda\tau}V^{-1}$ is given by:

$$e^{A\tau} = \frac{1}{\lambda_2 - \lambda_1} \begin{pmatrix} \lambda_2 e^{\lambda_1\tau} - \lambda_1 e^{\lambda_2\tau} & -e^{\lambda_1\tau} + e^{\lambda_2\tau} \\ \lambda_1\lambda_2\left(e^{\lambda_1\tau} - e^{\lambda_2\tau}\right) & -\lambda_1 e^{\lambda_1\tau} + \lambda_2 e^{\lambda_2\tau} \end{pmatrix} \;\;,$$

and it is useful to compute $\xi_T(u)$, defined as:

$$\xi_T(u) \;=\; \mathbb{E}\left[\sigma_u^2 | \mathcal{F}_T\right] \;=\; \mathbb{E}\left[p(X_u)^2 | \mathcal{F}_T\right] \;=\; \mathbb{E}\left[X_u^2 | \mathcal{F}_T\right] \;\;. \tag{4.14}$$

Carrying on the computation of the expected value in (4.14), we get to

$$\begin{aligned}\xi_T(u) &= \frac{1}{(\lambda_2-\lambda_1)^2}\left\{X_{1,T}\left(\lambda_2 e^{\lambda_1(u-T)}-\lambda_1 e^{\lambda_2(u-T)}\right)+X_{2,T}\left(e^{\lambda_2(u-T)}-e^{\lambda_1(u-T)}\right)\right\}^2 \\ &+ \frac{\gamma^2}{(\lambda_2-\lambda_1)^2}\left\{\frac{1}{2\lambda_2}\left(e^{2\lambda_2(u-T)}-1\right)-\frac{2}{\lambda_1+\lambda_2}\left(e^{(\lambda_1+\lambda_2)(u-T)}-1\right)+\frac{1}{2\lambda_1}\left(e^{2\lambda_1(u-T)}-1\right)\right\} \ .\end{aligned} \tag{4.15}$$

Plugging (4.15) into the VIX^2 expression in continuous time (2.13), calling $\sigma^2_{y_T}$ the variance of y_T, we get the following:

$$\begin{aligned}VIX_T^2 &= \frac{100^2}{\Delta}\int_T^{T+\Delta}\left\{\frac{1}{(\lambda_2-\lambda_1)^2}\left[X_{1,T}\left(\lambda_2 e^{\lambda_1(u-T))}-\lambda_1 e^{\lambda_2(u-T)}\right)+X_{2,T}\left(e^{\lambda_2(u-T)}-e^{\lambda_1(u-T)}\right)\right]^2+\sigma^2_{y_T}\right\}du \\ &= \frac{100^2}{\Delta}\cdot\frac{X_{1,T}^2}{(\lambda_2-\lambda_1)^2}\left\{\frac{\lambda_2^2}{2\lambda_1}\left(e^{2\lambda_1\Delta}-1\right)+\frac{\lambda_1^2}{2\lambda_2}\left(e^{2\lambda_2\Delta}-1\right)-\frac{2\lambda_1\lambda_2}{\lambda_1+\lambda_2}\left(e^{(\lambda_1+\lambda_2)\Delta}-1\right)\right\} \\ &+ \frac{100^2}{\Delta}\cdot\frac{X_{2,T}^2}{(\lambda_2-\lambda_1)^2}\left\{\frac{1}{2\lambda_1}\left(e^{2\lambda_1\Delta}-1\right)+\frac{1}{2\lambda_2}\left(e^{2\lambda_2\Delta}-1\right)-\frac{2}{\lambda_1+\lambda_2}\left(e^{(\lambda_1+\lambda_2)\Delta}-1\right)\right\} \\ &- \frac{100^2}{\Delta}\cdot\frac{2X_{1,T}X_{2,T}}{(\lambda_2-\lambda_1)^2}\left\{\frac{\lambda_2}{2\lambda_1}\left(e^{2\lambda_1\Delta}-1\right)+\frac{\lambda_1}{2\lambda_2}\left(e^{2\lambda_2\Delta}-1\right)-\left(e^{(\lambda_1+\lambda_2)\Delta}-1\right)\right\} \\ &+ \frac{100^2}{\Delta}\cdot\frac{\gamma^2}{(\lambda_2-\lambda_1)^2}\left\{\Delta\left(\frac{2}{\lambda_1+\lambda_2}-\frac{1}{2\lambda_2}-\frac{1}{2\lambda_1}\right)\right\} \\ &+ \frac{100^2}{\Delta}\cdot\frac{\gamma^2}{(\lambda_2-\lambda_1)^2}\left\{\frac{1}{4\lambda_2^2}\left(e^{2\lambda_2\Delta}-1\right)-\frac{2}{(\lambda_1+\lambda_2)^2}\left(e^{(\lambda_1+\lambda_2)\Delta}-1\right)+\frac{1}{4\lambda_1^2}\left(e^{2\lambda_1\Delta}-1\right)\right\} \ .\end{aligned}$$

The state vector X_t follows a Gaussian distribution, with mean m_{X_t} and variance

$$\mathrm{Var}(X_t) := \Sigma_t = \gamma^2\int_0^t e^{Au}e_2e_2^Te^{A^Tu}du \ .$$

Explicitly, the covariance matrix Σ_t is:

$$\Sigma_t = \frac{\gamma^2}{(\lambda_2-\lambda_1)^2}\begin{pmatrix}\frac{1}{2\lambda_1}\left(e^{2\lambda_1 t}-1\right)+\frac{1}{2\lambda_2}\left(e^{2\lambda_2 t}-1\right)-\frac{2}{\lambda_1+\lambda_2}\left(e^{(\lambda_1+\lambda_2)t}-1\right) & \frac{1}{2}\left(e^{2\lambda_1 t}-1\right)+\frac{1}{2}\left(e^{2\lambda_2 t}-1\right)-e^{(\lambda_1+\lambda_2)t}+1 \\ \frac{1}{2}\left(e^{2\lambda_1 t}-1\right)+\frac{1}{2}\left(e^{2\lambda_2 t}-1\right)-e^{(\lambda_1+\lambda_2)t}+1 & \frac{\lambda_1}{2}\left(e^{2\lambda_1 t}-1\right)+\frac{\lambda_2}{2}\left(e^{2\lambda_2 t}-1\right)-\frac{2\lambda_1\lambda_2}{\lambda_1+\lambda_2}\left(e^{(\lambda_1+\lambda_2)t}-1\right)\end{pmatrix} \ .$$

Finally, pricing a VIX derivative with payoff function V is achieved by integrating against the two-dimensional Gaussian density:

$$\mathbb{E}\left[V\left(\sqrt{\mathrm{VIX}_T^2}\right)\right] = \frac{1}{2\pi\sqrt{\det(\Sigma_T)}}\int_{\mathbb{R}^2}V\left(\sqrt{h(x)}\right)e^{-\frac{1}{2}(x-mX_T)^T\Sigma_T^{-1}(x-m_{X_T})}dx \ , \tag{4.16}$$

where $h(\cdot)$ represents the polynomial expression of $\text{VIX}_T^2(\cdot)$, and m_{X_T} and Σ_T denote the mean vector and covariance matrix, respectively.

4.5. Summary

In this Chapter, we introduced the CAR(2) stochastic volatility model and explored its application to the joint calibration of SPX and VIX derivatives. We derived new explicit pricing formulae for VIX futures and options by computing the forward variance as a polynomial function of the Gaussian state variables. These closed-form expressions allow for efficient, semi-analytical pricing of VIX derivatives without the need of Monte Carlo simulations. For SPX options, the model's affine structure enabled the derivation of a system of ordinary differential equations for the characteristic function of the log-price process, facilitating pricing via FFT-based techniques. The CAR(2) model successfully captures essential market features—such as the VIX futures term structure—while maintaining a parsimonious set of parameters and computational tractability. These qualities suggest that CAR(2) can serve as a promising framework for volatility modeling in joint calibration settings, particularly when balancing model expressiveness with tractability.

5 | Calibration

In this Chapter, we describe the dataset and the methodology employed for calibrating the stochastic volatility models under consideration. Particular emphasis is placed on the calibration cascade technique, a structured approach that sequentially aligns key market quantities—such as discount factors and forward prices—before calibrating to SPX and VIX options and futures. This methodology improves both the stability and interpretability of the joint calibration process.

5.1. Dataset

We evaluate our models using two closely aligned datasets from June 20, 2021, and July 20, 2021, containing detailed data from the SPX and VIX options markets. To complement these, we include a supplementary dataset from June 5, 2025, which provides directly observed market levels for VIX futures. This allows us to test the robustness of our joint calibration framework not only across different market periods but also against actual VIX futures data, rather than relying on synthetic estimates derived from options. Each dataset includes monthly maturities for both call and put options on the aforementioned indices. Since the maturities of the VIX options do not precisely coincide with the SPX options, there are minor differences, usually limited to a few days. This small discrepancy in maturities is considered negligible for our analysis.

This offset in maturities primarily comes from the expiration schedules of these instruments. VIX options typically expire on Wednesdays, while SPX options usually expire on Fridays. These conventions are defined by Cboe, as only SPX options with Friday expirations are used to calculate the VIX Index. [1]

For SPX options specifically, the study focuses only on options within the 25% delta range, as we want to concentrate the calibration on options with sufficient liquidity and meaningful sensitivity to underlying asset price changes. This choice provides a balanced

[1] Cboe Global Markets:VIX options and SPX options product specification (2025). Available at: https://www.cboe.com/tradable_products/vix/vix_options/specifications/ https://www.cboe.com/tradable_products/sp_500/spx_options/specifications/

representation of market expectations without extreme in-the-money or out-of-the-money biases.

Employing datasets from two adjacent months allows us to assess the temporal consistency and robustness of the models.

5.1.1. VIX and SPX implicit interest rates

In this Section we extract the discount factors implied by market option prices, following Azzone and Baviera (2021). These discount factors reflect the market-implicit interest rates which, when compared to the USD Overnight Indexed Swap (OIS) curve, allow us to analyze the presence of any risk or liquidity premiums.

For every available maturity t, we take the set of N strikes K_i and corresponding put P_i and call C_i option prices.

$$B(t_0,t) = -\frac{\sum_{i=1}^{N}\left(K_i-\hat{K}\right)\left(\mathcal{G}_i-\hat{\mathcal{G}}\right)}{\sum_{i=1}^{N}\left(K_i-\hat{K}\right)^2} \quad , \tag{5.1}$$

where

$$\mathcal{G}_i = C_i - P_i \quad , \qquad \hat{\mathcal{G}} = \frac{1}{N}\sum_{i=1}^{N}\mathcal{G}_i \quad , \qquad \hat{K} = \frac{1}{N}\sum_{i=1}^{N}K_i \quad . \tag{5.2}$$

And, for any K_i, we get the price of the forward contract as

$$F(t_0,t) = \frac{\mathcal{G}_i}{B(t_0,t)} + K_i \quad . \tag{5.3}$$

The funding spread can be defined as

$$s := \frac{1}{t-t_0}\log\left(\frac{B(t_0,t)}{\overline{B}(t_0,t)}\right) \quad , \tag{5.4}$$

where $\overline{B}(t_0,t)$ is obtained from the bootstrap of the OIS curve.

Overall, both markets exhibit a positive spread over the OIS curve, resulting in a cost of funding of 45 basis points for VIX and 42 for SPX. In both cases, the spread shows higher variance for shorter maturities, before stabilizing around the average value, as we can see in Figure 5.1.

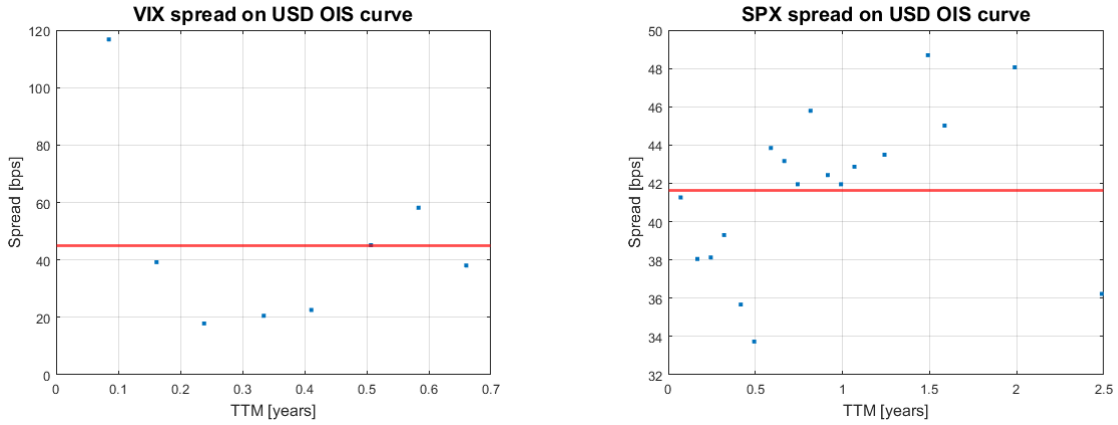

(a) Spread over OIS, VIX options.

(b) Spread over OIS, SPX options.

Figure 5.1: Spread over OIS in the VIX and SPX options market (June 20, 2021 data). Spread plotted against time-to-maturity. The red line is the average spread.

5.2. Methodology

In this Section we explain the multiple-steps calibration cascade method, where each step builds on the results of the previous one to ensure consistency and model coherence.
The process begins with the exact calibration of the forward levels and discount factors for both markets. These results are the foundation of the calibration procedure.
The second step starts with fitting the model to VIX futures and VIX option smiles, ensuring the volatility structure is well captured.
Then, based on the results of the VIX calibration, we finalize the process by calibrating the remaining parameters to fit the SPX volatility smile.

5.2.1. Exact calibration of Forward and Discount factors

In the first step, we perform an exact calibration of the discount factors and forward values.
Market-implied discount factors are computed using (5.1), as explained in the previous Section. We have slightly different discount curves for VIX and SPX options.
Additionally, the forward levels are extracted as in (5.3) to serve as a baseline for subsequent volatility surface calibration.

5.2.2. Joint calibration of VIX forward and VIX-SPX smiles

The second step of the calibration cascade begins with the VIX market, incorporating both options and their corresponding futures.
This approach is motivated not only by the generally greater liquidity of VIX futures compared to options, but also by their essential role in implied volatility analysis. Although our models do not rely on the Black formula for pricing, we use this formula to extract implied volatilities, and it requires the VIX future as the underlying input. Therefore, an accurate calibration of VIX futures is crucial to ensure the consistency of the implied volatility surface and its alignment with observed market data.
To calibrate the model parameters θ, for each maturity slice, we minimize the following loss function:

$$J_{VIX}(\theta) = w_{fut}\sqrt{\left(F^{\theta}_{VIX} - F^{mkt}_{VIX}\right)^2} + w_{opt}\sqrt{\sum_j \left(\sigma^{\theta}_{VIX}(K_j) - \sigma^{mkt}_{VIX}(K_j)\right)^2} \ , \tag{5.5}$$

where $\sigma^{\theta}_{VIX}(K_j)$ and $\sigma^{mkt}_{VIX}(K_j)$ are model and market implied volatility values for the maturity slice, with strike K_j. F^{θ}_{VIX} and F^{mkt}_{VIX} are model and market futures price with the same maturity as the options. We choose w_{fut} to be greater than w_{opt} for our analysis.

Once the VIX derivatives are successfully calibrated, we can finally proceed to the final step of the cascade: the calibration of SPX options. All parameters related to the interest rate curve and the volatility dynamics inferred from the VIX market are fixed, drastically reducing the dimensionality of the problem, leaving only two parameters to calibrate (σ, ρ).
Not only does this focused calibration simplify the optimization problem, enhancing performance, but it also ensures that the SPX smile is consistent with both the observed option prices and the VIX market-implied volatility structure.
To obtain the price of an SPX call option, we use the Lewis (2001) formula.
Let x be the log-moneyness, $\phi_{t_0,t}(\xi)$ be analytic in the horizontal strip $\{\xi \in \mathbb{C} : -1 \leq \mathrm{Im}(\xi) \leq 0\}$, the call option price $c(x)$ is

$$\frac{c(x)}{B(t_0,t)F(t_0,t)} = 1 - \frac{e^{-x/2}}{2\pi}\int_{-\infty}^{\infty} e^{-ixp}\phi_{t_0,t}\Big(-u - \frac{i}{2}\Big)\frac{1}{u^2 + \frac{1}{4}}\,du \ . \tag{5.6}$$

To calibrate the remaining parameters, for each maturity slice, we minimize

$$J_{SPX}(\theta) = \frac{\sum_j \left(c^{\theta}_{SPX}(K_j) - c^{mkt}_{SPX}(K_j)\right)^2}{c^{mkt}_{ask}(K_j) - c^{mkt}_{bid}(K_j)} \ . \tag{5.7}$$

where $c^{\theta}_{SPX}(K_j)$ and $c^{mkt}_{SPX}(K_j)$ are model and market call option prices for the maturity slice, with strike K_j, and $c^{mkt}_{ask}(K_j) - c^{mkt}_{bid}(K_j)$ is the market bid-ask spread.

The calibration is performed independently on each maturity slice, allowing flexibility across the term structure. For the VIX market, we calibrate on the observed maturity T. For the SPX options, we focus on the corresponding maturity $T_{SPX} = T + \Delta$, with $\Delta = 30$ days. While some works in the literature consider both SPX maturities T and $T + \Delta$ (see Jaber et al. (2023), Cuchiero et al. (2024)), we restrict our analysis to $T + \Delta$ SPX maturity only.
This choice is motivated by the definition of the VIX itself, which measures the market's expected SPX implied volatility over the following 30 days. By including the T maturity for SPX options, we would risk to introduce a temporal mismatch, as those options expire before the start of the volatility window implied by the VIX index. By focusing only on the $T + \Delta$ maturity, we ensure consistency between the two markets, which is crucial for a coherent joint calibration.

5.3. Summary

In this Chapter we have proposed a calibration cascade method to jointly calibrate VIX and SPX markets. This approach offers several advantages: it ensures consistency between the two markets, reduces the dimensionality of the overall minimization problem by decomposing it into smaller, more manageable stages, and enhances interpretability by finding market-implied quantities (such as discount factors and forward levels), before fine-tuning the SPX volatility smile using only a minimal set of remaining parameters.

6 | Results

In this Chapter, we present and compare the VIX and SPX implied volatility smiles produced by the three models, calibrated following the methodology outlined in Chapter 5. For each model, we provide visual comparisons and discuss the main strengths and limitations observed in the joint calibration results. Calibrated parameter values are also reported. In addition to the original 2021 datasets, we include selected results from the June 2025 dataset, which includes observed VIX futures levels, offering a more direct benchmark for evaluating the accuracy of the VIX leg of the models. Additional calibration results for the July 20, 2021 dataset are provided in Appendix B.

6.1. Heston

The Heston model is able to match the observed VIX futures level across maturities with good precision, as illustrated in Figures 6.1, 6.2 and 6.3. However, it struggles to capture the shape of the VIX implied volatility smile, failing to reproduce the upward-sloping and convex structure observed in market data. As a result, the overall quality of the VIX option calibration is poor, despite matching the future level itself.
This shortcoming propagates to the SPX leg of the calibration. Since the calibration procedure is sequential, starting from the VIX market, the model enters the SPX calibration stage with suboptimal parameter values. Consequently, the model also underperforms on the SPX side: while it may align with the general slope of the implied volatility skew, it fails to reproduce the correct smile curvature and market levels consistently across strikes.

Nevertheless, one structural feature of the Heston model remains consistent with empirical literature: the need for a strongly negative correlation between the log-return and volatility processes. As shown in Table 6.1, the calibrated correlation values typically collapse to $\rho \leq -0.9$, reflecting the leverage effect commonly observed in equity markets. Although this value might appear extreme, it is supported by results in Heston and Nandi (1998), where a GARCH-type model yields $\rho = -1$ under maximum likelihood estimation.

To further evaluate the model's performance under more realistic conditions, we include

calibration results from the June 5, 2025 dataset, shown in Figure 6.4, where actual VIX futures quotes are available. These observed futures levels provide a stricter benchmark than synthetic values extracted from options data. In this case, the Heston model struggles not only to fit the implied volatility smile, but also to accurately reproduce the VIX futures levels across maturities. The model-implied futures tend to fall below market levels, and the pricing of VIX options remains poor. These discrepancies confirm that the model's limitations are structural and persist even when calibration is performed on real market data. The inability to reconcile both the level and the shape of the VIX derivatives limits the effectiveness of the Heston model in joint calibration tasks and raises concerns about its suitability for accurately capturing the joint dynamics of SPX and VIX markets.

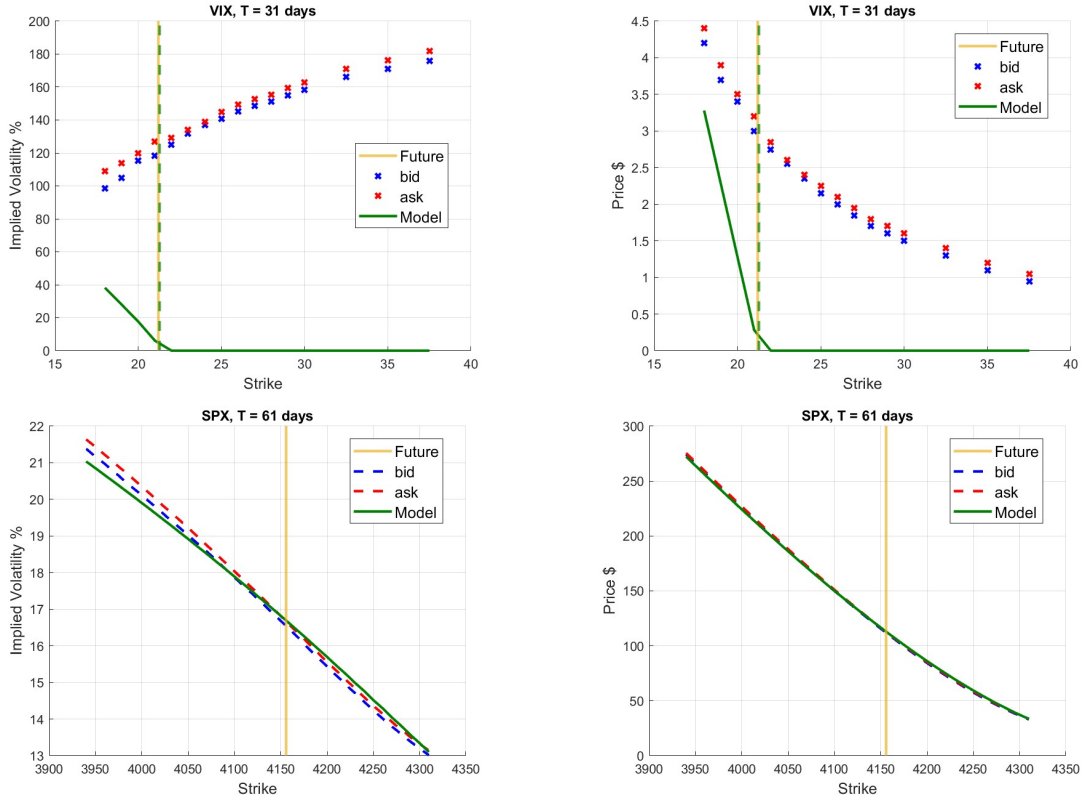

Figure 6.1: Maturity slice T_1: SPX-VIX joint calibration results of Heston model on June 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

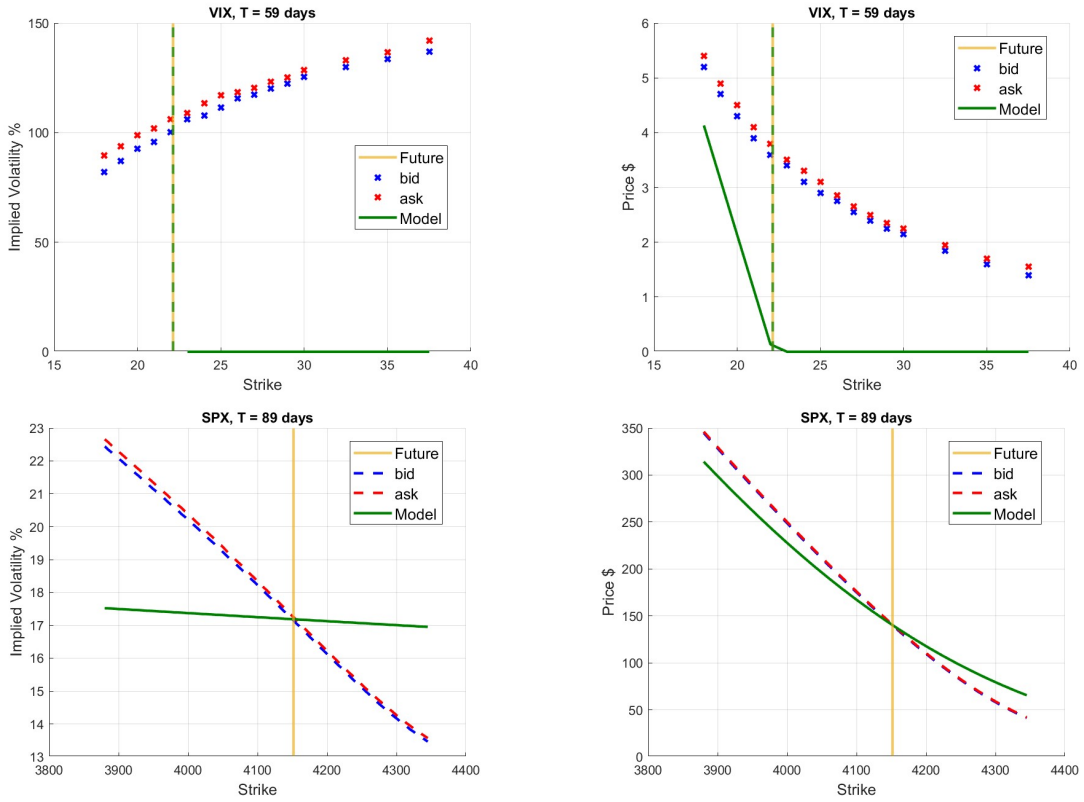

Figure 6.2: Maturity slice T_2: SPX-VIX joint calibration results of Heston model on June 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

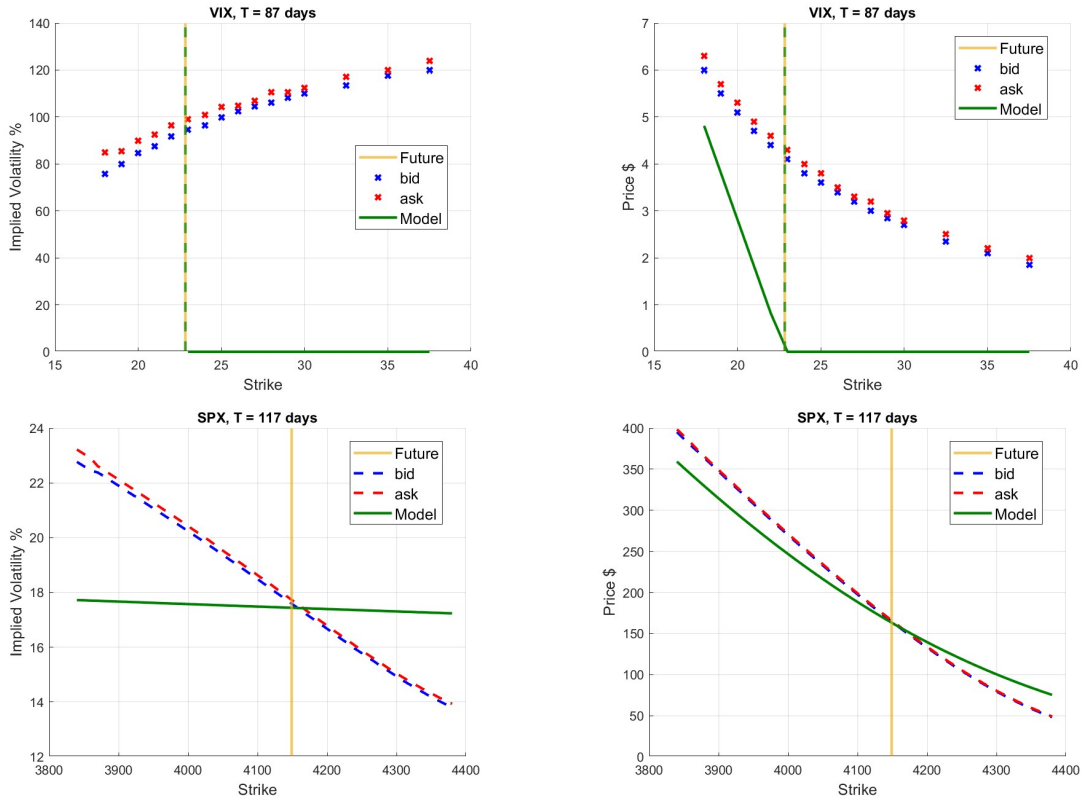

Figure 6.3: Maturity slice T_3: SPX-VIX joint calibration results of Heston model on June 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

Heston model calibrated parameters

	σ	γ	ρ	a	x_0
June 20, 2021 dataset					
$T_{\text{VIX}} = 31\text{d}$	0.6516	1.7234	-0.9418	1.1210	0.0000
$T_{\text{VIX}} = 59\text{d}$	0.7602	1.7902	-0.9023	0.0292	0.0478
$T_{\text{VIX}} = 87\text{d}$	0.7493	1.8072	-0.9356	0.0195	0.0523
June 5, 2025 dataset					
$T_{\text{VIX}} = 13\text{d}$	0.7788	1.7890	-0.8746	0.7392	0.0046
$T_{\text{VIX}} = 41\text{d}$	0.6605	1.7911	-0.9000	0.3991	0.0290
$T_{\text{VIX}} = 76\text{d}$	0.5588	1.7559	-0.9000	0.6183	0.0210

Table 6.1: Heston model VIX-SPX jointly calibrated parameters for the June 2021 and June 2025 datasets.

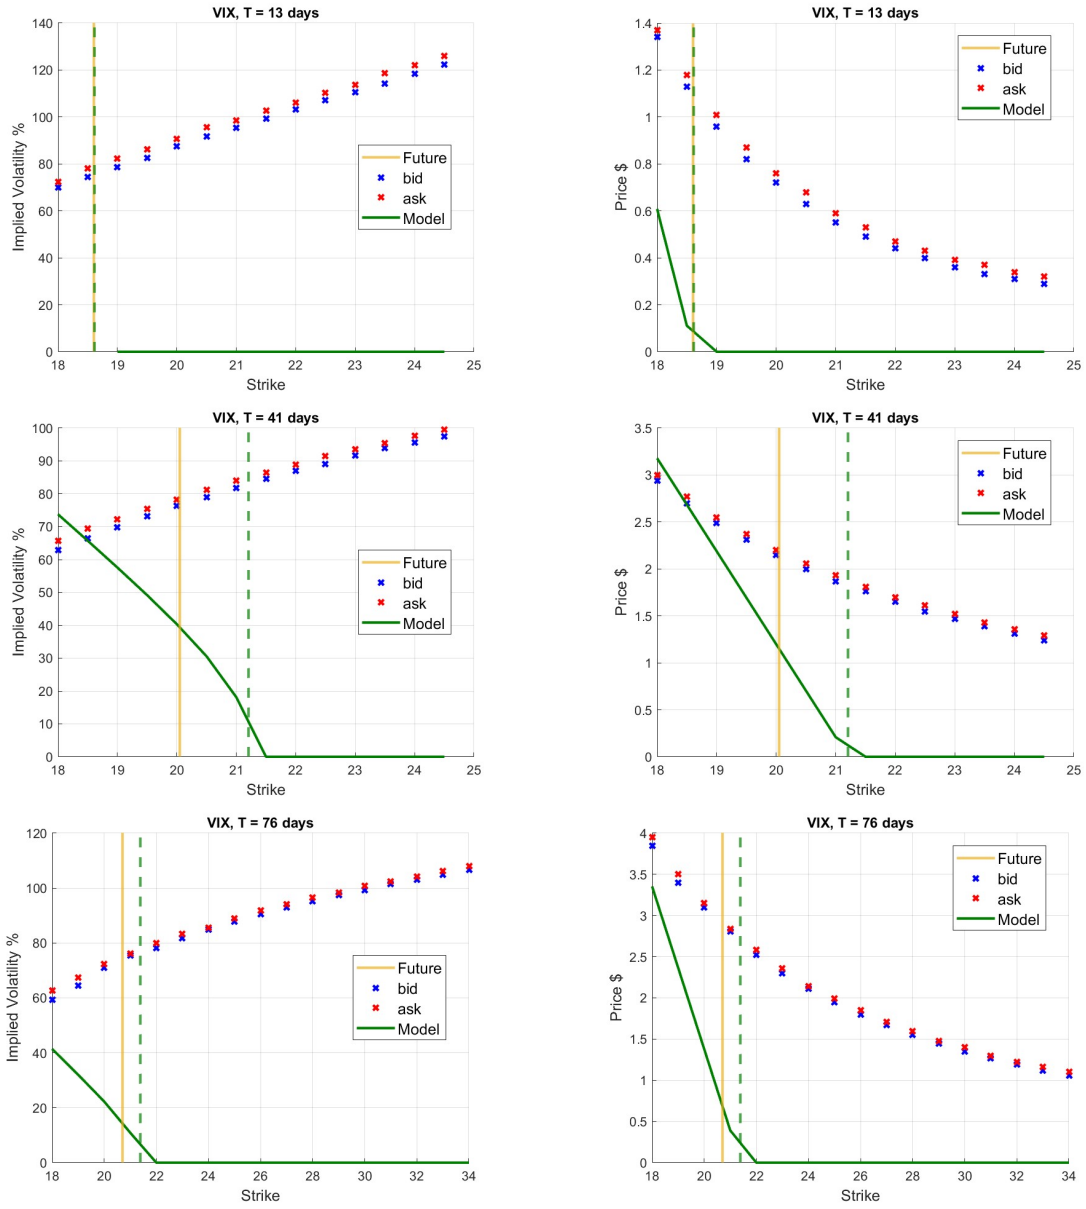

Figure 6.4: Maturity slices T_1, T_2 and T_3: VIX calibration results of the Heston model on June 5, 2025 data. Each row corresponds to a separate calibration for a distinct maturity. Left column: market implied volatilities (bid/ask in blue/red) compared to the model output (green). Right column: corresponding VIX option prices. The vertical gold line marks the observed VIX future level; the vertical green dashed line indicates the model-implied VIX future level. Calibration is performed independently for each maturity slice, allowing for maturity-specific parameter estimation.

6.2. Ornstein–Uhlenbeck

The OU model captures the general upward-sloping shape of the VIX implied volatility smile, but it overestimates volatility at lower strikes, while strongly underestimating it at higher strikes, particularly for far in-the-money call options. As shown in Figures 6.5, 6.6 and 6.7, this mismatch becomes more pronounced as the strike moves away from the VIX future level (indicated by the gold vertical line), where the model curve (green line) flattens prematurely. As a result, the model fails to reproduce the full convexity observed in market data. Nevertheless, in the top-right panel of Figure 6.5, the model performs reasonably well in reproducing VIX option prices. This is largely due to its ability to match the market future level almost perfectly, which serves as the main driver of VIX option valuation. Therefore, even though the implied volatility fit is suboptimal, the model achieves acceptable pricing accuracy.
Regarding SPX options, the model's performance improves as time to maturity increases. In Figure 6.5 the model overestimates implied volatility across all strikes, resulting in higher-than-observed option prices. However, in Figures 6.6 and 6.7, the fit improves considerably. The model matches implied volatilities almost perfectly around the at-the-money region, leading to an excellent alignment with market prices.
As observed previously in Section 6.1, the calibrated correlation parameter remains strongly negative, with $\rho \approx -0.95$. This reflects the trade-off between using lower ρ values to steepen the smile and increasing σ (vol-of-vol) to vertically shift the curve upward (Table 6.2).

To further evaluate the model under realistic market conditions, we calibrate it to the June 5, 2025 dataset, which includes directly observed VIX futures prices (see Figure 6.8). Interestingly, the model performs slightly better in this setting: the fit to the VIX futures term structure is almost exact across maturities, and the pricing of VIX options remains consistent with market levels. While the implied volatility smile is still a bit flattened compared to observed convexity—particularly at higher strikes—the accuracy in matching futures levels contributes to acceptable option pricing. These results suggest that the OU model benefits from anchoring to actual futures data, yielding more reliable outcomes than when relying on synthetic estimates derived from options.

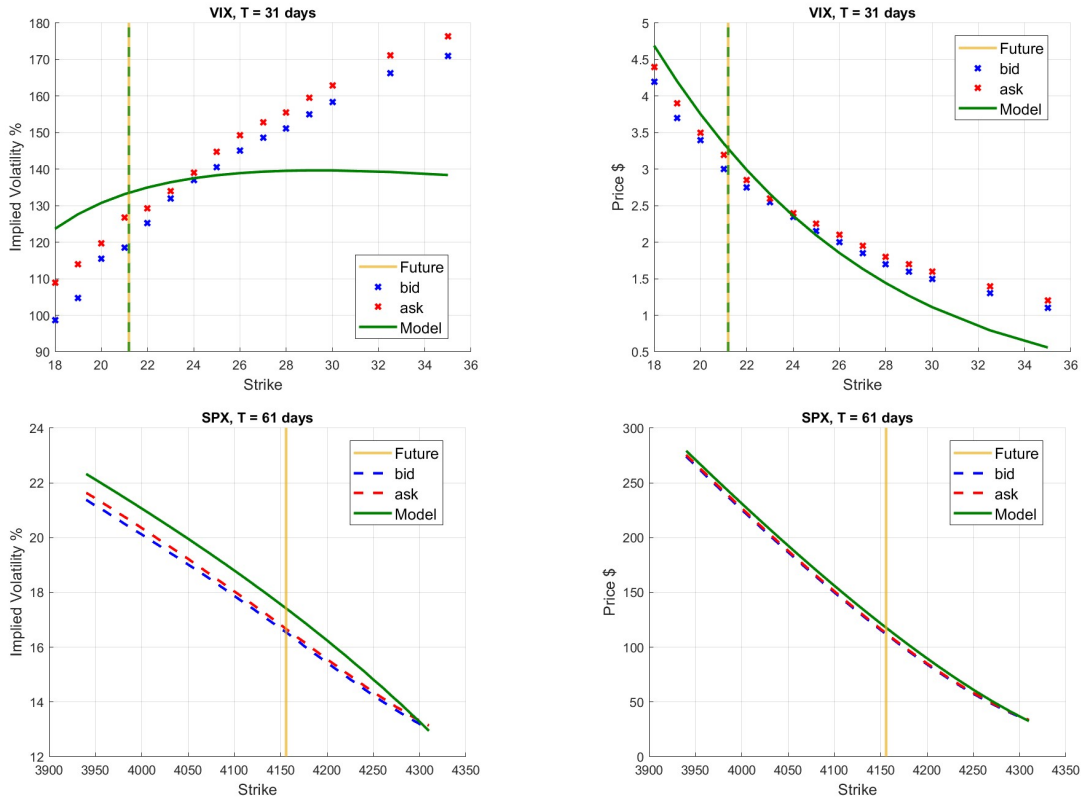

Figure 6.5: Maturity slice T_1: SPX-VIX joint calibration results of OU model on June 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

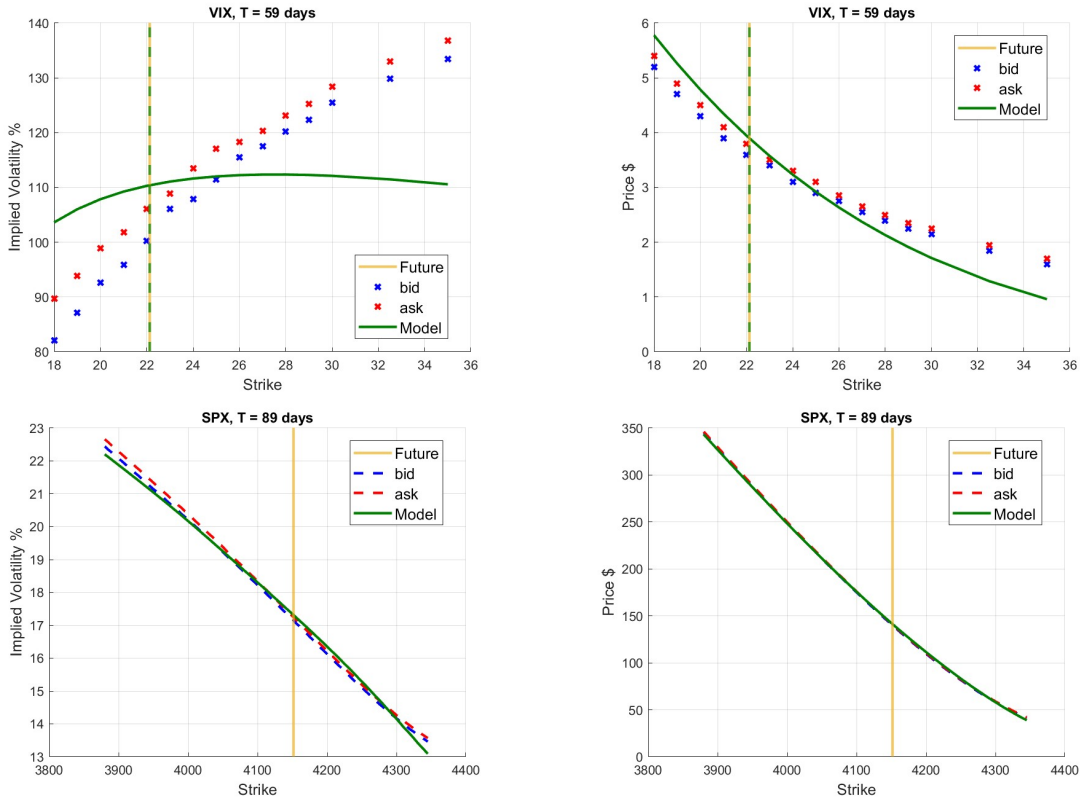

Figure 6.6: Maturity slice T_2: SPX-VIX joint calibration results of OU model on June 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

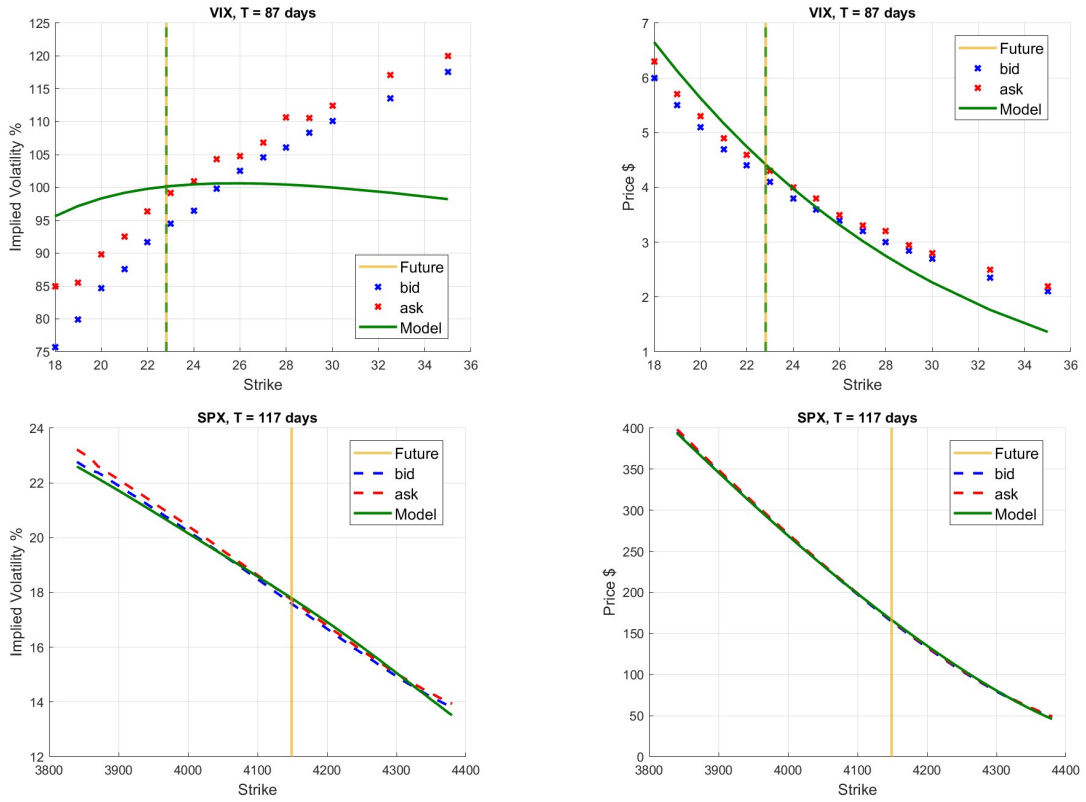

Figure 6.7: Maturity slice T_3: SPX-VIX joint calibration results of OU model on June 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

OU model calibrated parameters

	μ	σ	γ	ρ	a	x_0
June 20, 2021 dataset						
$T_{\text{VIX}} = 31$d	0.1289	1.0817	0.6927	-0.9680	1.8501	0.0469
$T_{\text{VIX}} = 59$d	0.0789	0.9915	0.6724	-0.9511	3.0785	0.0852
$T_{\text{VIX}} = 87$d	0.0773	0.9604	0.6215	-0.9435	2.6455	0.1078
June 5, 2025 dataset						
$T_{\text{VIX}} = 13$d	0.0286	0.9604	1.1793	-0.7935	18.9784	0.0842
$T_{\text{VIX}} = 41$d	0.0330	0.9210	0.9382	-0.8880	10.4124	0.1516
$T_{\text{VIX}} = 76$d	0.0488	0.8366	0.7376	-0.9082	6.1082	0.2186

Table 6.2: OU model VIX-SPX jointly calibrated parameters for the June 2021 and June 2025 datasets.

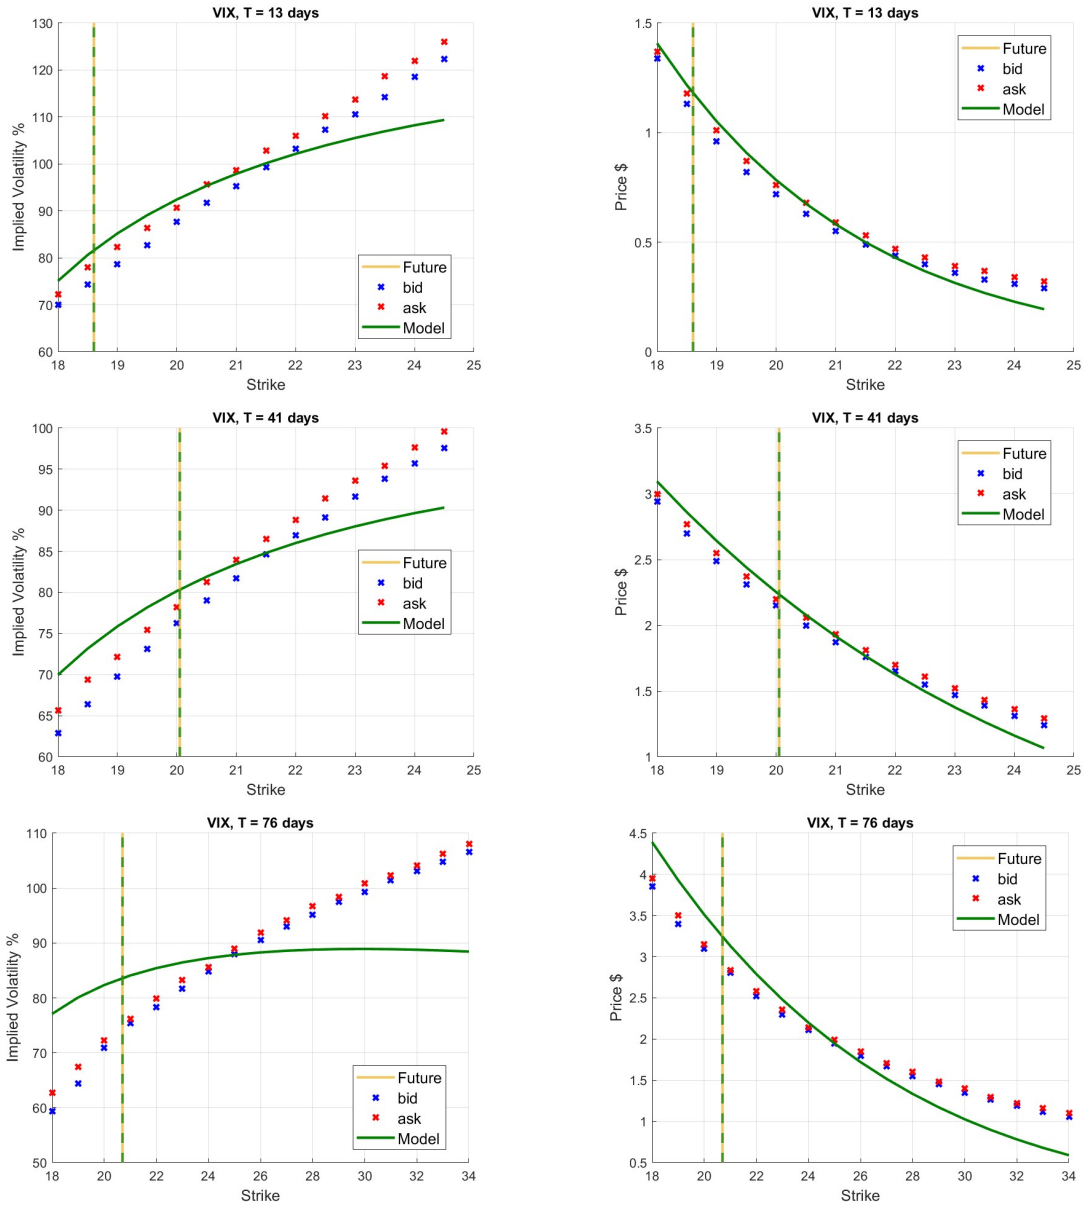

Figure 6.8: Maturity slices T_1, T_2 and T_3: VIX calibration results of the OU model on June 5, 2025 data. Each row corresponds to a separate calibration for a distinct maturity. Left column: market implied volatilities (bid/ask in blue/red) compared to the model output (green). Right column: corresponding VIX option prices. The vertical gold line marks the observed VIX future level; the vertical green dashed line indicates the model-implied VIX future level. Calibration is performed independently for each maturity slice, allowing for maturity-specific parameter estimation.

6.3. Continuous Autoregressive

One of the key successes of the CAR(2) model lies in its ability to accurately match the VIX future level across different maturities, as shown in Figures 6.9, 6.10 and 6.11. This represents a significant improvement over the Heston model (Figure 6.1), which fails to capture the VIX future level. Given the VIX future's high liquidity and its central role in VIX option pricing, this match enhances the practical relevance and robustness of the CAR(2) calibration. In addition, the model achieves a reasonable fit for VIX option prices, particularly around the at-the-money region (top-right panel), even though the full implied volatility smile (top-left panel) is not perfectly captured.
For SPX options, the model shows a good qualitative match in both implied volatilities and option prices, especially around the at-the-money region. While option prices are calibrated quite well, the model struggles to reproduce the steep implied volatility skew. As observed in Sections 6.1 and 6.2, the correlation parameter collapses to $\rho = -1$, highlighting a structural limitation also in this model's ability to decouple the effects of correlation and vol-of-vol. In particular, the model struggles to simultaneously reproduce the steepness of the implied volatility skew and the overall level of the volatility smile, due to the coupled effects of ρ and σ. Lowering ρ steepens the smile, while increasing σ shifts it upward, but the model does not allow sufficient flexibility to control these effects independently.
Finally, we observe that the calibrated parameters remain relatively stable across maturities, with the exception of the mean-reverting coefficients α_1 and α_2, which tend to increase with the time to maturity, as we can see in Table 6.3.

We further assess the CAR(2) model on the June 5, 2025 dataset, which includes observed VIX futures prices (see Figure 6.12). As in the 2021 dataset, the model reproduces the VIX futures level well for shorter maturities. However, at the longest maturity ($T = 76$ days), the model underestimates the VIX future level. Despite this, the pricing of VIX options remains acceptable across maturities, particularly near the at-the-money region. As observed previously, the model continues to struggle with replicating the steepness and convexity of the implied volatility smile—especially at short maturities—flattening the smile too early and missing the skew visible in market data. These results highlight the model's robustness in short-term VIX futures fitting, but also its structural limitations in capturing the full volatility surface as maturity increases.

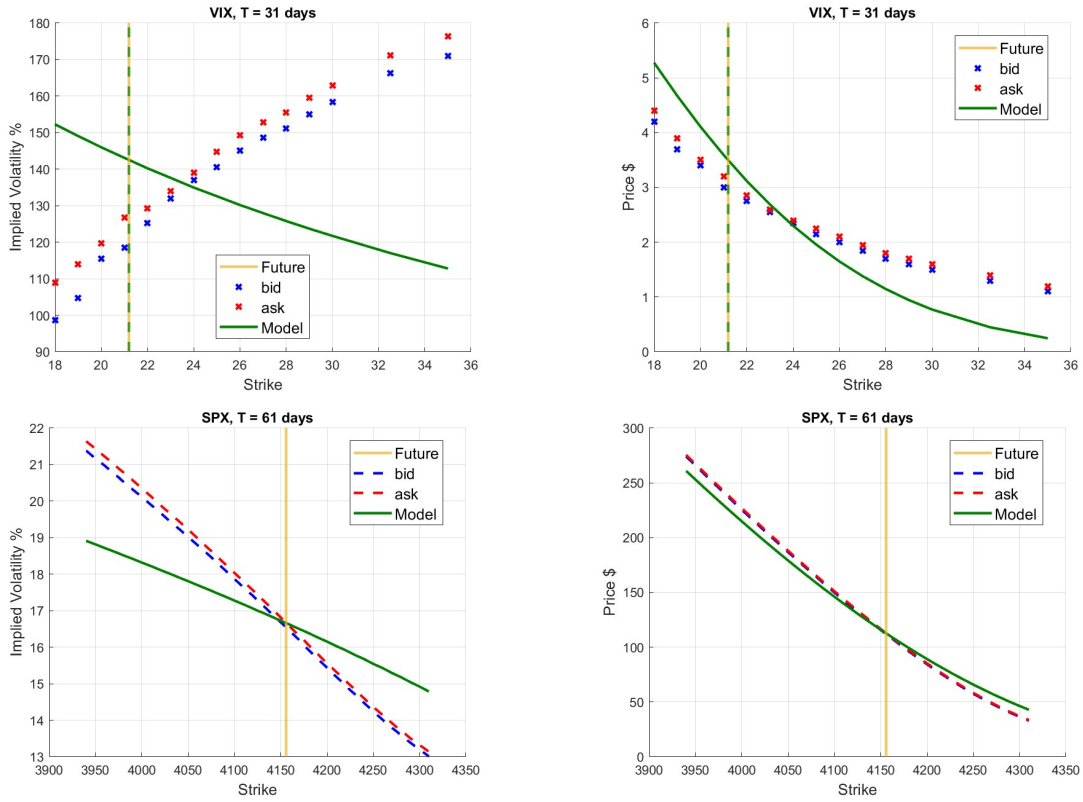

Figure 6.9: Maturity slice T_1: SPX-VIX joint calibration results of CAR(2) model on June 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

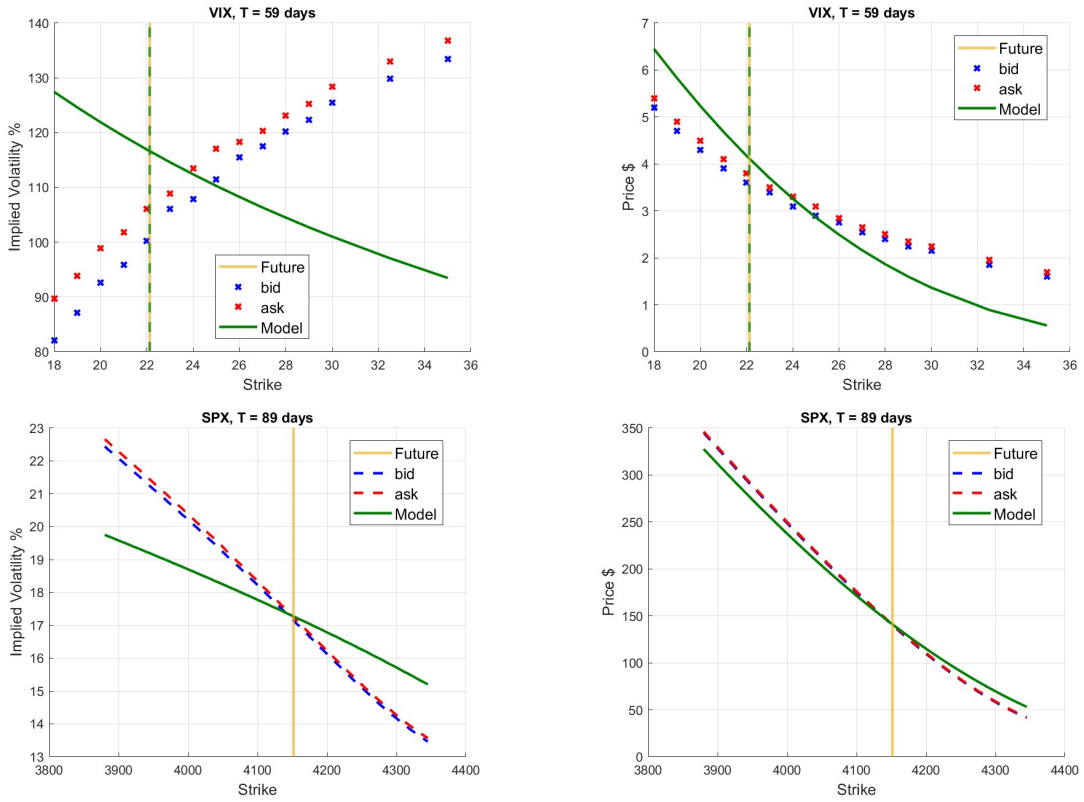

Figure 6.10: Maturity slice T_2: SPX-VIX joint calibration results of CAR(2) model on June 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

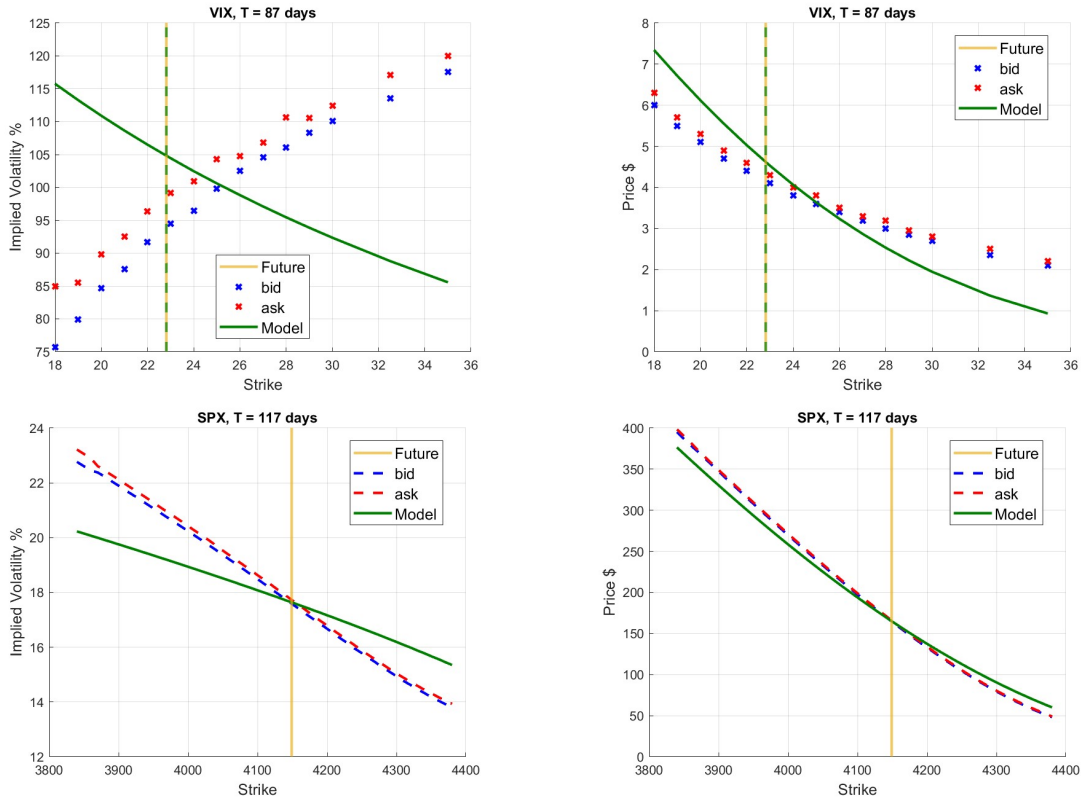

Figure 6.11: Maturity slice T_3: SPX-VIX joint calibration results of CAR(2) model on June 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

CAR(2) model calibrated parameters

	σ	γ	ρ	α_1	α_2	x_1	x_2
June 20, 2021 dataset							
$T_{\text{VIX}} = 31$d	0.8389	3.6598	-1.0000	0.6843	0.0616	0.1629	0.3671
$T_{\text{VIX}} = 59$d	0.8553	2.9073	-1.0000	4.9499	2.3885	0.1485	0.5935
$T_{\text{VIX}} = 87$d	0.8400	2.3059	-1.0000	5.0122	2.5185	0.1547	0.5425
June 5, 2025 dataset							
$T_{\text{VIX}} = 13$d	0.8362	8.8000	-1.0000	39.6337	32.8034	0.1769	0.6354
$T_{\text{VIX}} = 41$d	1.0232	3.8485	-1.0000	14.4572	38.3223	0.0763	-4.8844
$T_{\text{VIX}} = 76$d	1.2059	29.6174	-0.8408	59.8708	59.7653	0.0982	-5.8684

Table 6.3: CAR(2) model VIX-SPX jointly calibrated parameters for the June 2021 and June 2025 datasets.

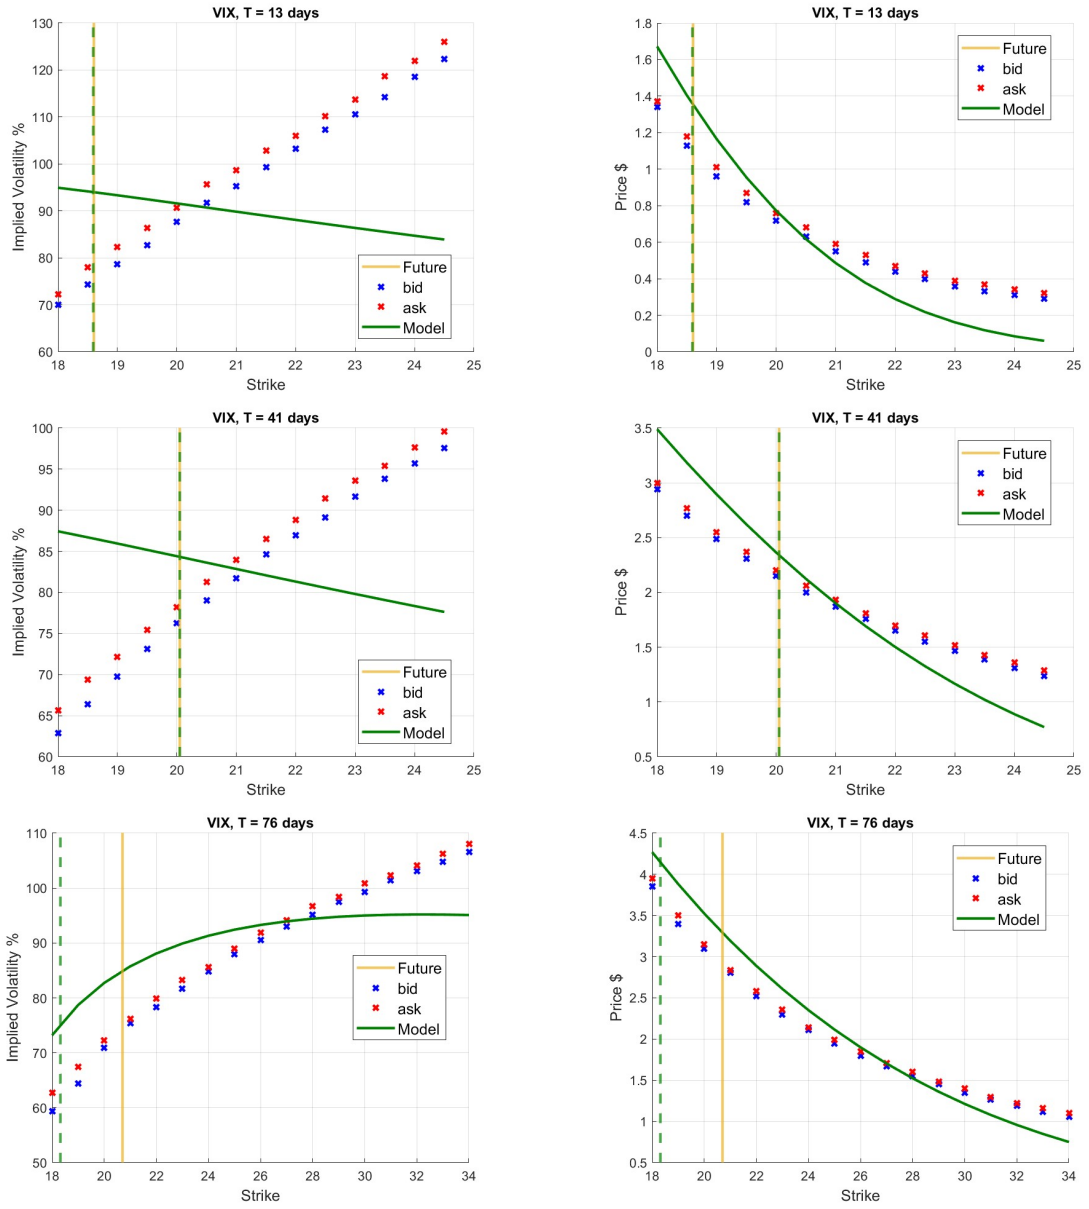

Figure 6.12: Maturity slices T_1, T_2 and T_3: VIX calibration results of the CAR(2) model on June 5, 2025 data. Each row corresponds to a separate calibration for a distinct maturity. Left column: market implied volatilities (bid/ask in blue/red) compared to the model output (green). Right column: corresponding VIX option prices. The vertical gold line marks the observed VIX future level; the vertical green dashed line indicates the model-implied VIX future level. Calibration is performed independently for each maturity slice, allowing for maturity-specific parameter estimation.

6.4. Summary

This Chapter has presented and analyzed the joint calibration results of three continuous time stochastic volatility models—Heston, Ornstein-Uhlenbeck, and the proposed CAR(2) model—applied to SPX and VIX derivatives.

The Heston model exhibited overall poor performance in the joint calibration task. Although the steep SPX implied volatility skew was captured through a highly negative correlation parameter ρ, consistent with the leverage effect observed in equity markets, the model failed to accurately price SPX options across strikes and maturities. Its performance on the VIX leg was also unsatisfactory: the model struggled to reproduce the implied volatility smile and could not match the observed VIX futures levels, even when real market quotes were used. These shortcomings highlight structural limitations that prevent the Heston model from effectively reconciling the joint behavior of SPX and VIX derivatives.

The OU model achieved stable calibration and acceptable option price fits, especially at longer maturities. Although it struggled to capture the full convexity of the VIX implied volatility smile, it provided a very accurate match of the SPX one, especially around the at-the-money region, consistently staying within the market bid-ask spread. This suggests that despite its structural simplicity, the OU model remains a good candidate for the joint calibration problem, offering a balanced trade-off between tractability and accuracy.

The CAR(2) model, introduced in this thesis as part of the joint calibration framework, demonstrated both strengths and limitations compared to the benchmark models. It consistently matched the VIX future level across maturities—a notable improvement on the Heston model, particularly given the future's central role in option pricing. The model also achieved overall good pricing accuracy for both SPX and VIX options. However, like the other models, it required a correlation parameter that collapsed to $\rho = -1$, reflecting the unresolved trade-off between steepening the smile and raising its overall level due to the interdependent roles of correlation and vol-of-vol parameters, which restrict the model's ability to control skew and level independently.

An important feature shared by all three models studied in this work—Heston, OU and CAR(2)—is their relatively low dimensionality, each involving only five to seven parameters. This makes them appealing for practical calibration tasks, especially when compared to more complex models requiring higher-dimensional optimization.
Moreover, the calibration procedure is Monte Carlo-free, relying instead on semi-analytical pricing techniques. This significantly reduces computational cost and improves numeri-

cal stability, further enhancing the models' potential applicability in real-world settings, where efficiency and robustness are essential.

7 | Conclusions

This thesis addressed the challenge of jointly calibrating SPX and VIX derivatives using continuous-time stochastic volatility models. In particular, we utilized the Heston and OU models, and the newly introduced CAR(2) model for the joint calibration framework. Through a multi-stage calibration cascade, we aligned market-implied quantities—such as discount factors and forwards—and VIX volatility, before fitting the SPX smile, enhancing both the interpretability and the coherence of the joint calibration process.

In the following, we summarize the main contributions of this thesis.

- In this work, we analyzed and implemented the Heston and OU models, evaluating their capacity to price VIX derivatives within a joint calibration framework. In doing so, we extended the theoretical insights of Mechkov (2015) by illustrating how the fast mean-reversion limit of the Heston model connects to the Normal Inverse Gaussian distribution. Specifically we provided an explicit computation of the characteristic function and derived a direct mapping between the parameters of the two models.
- This thesis introduced a new application of the second-order Continuous Autoregressive (CAR(2)) process to equity volatility modeling. Exploiting its continuous and affine structure, we derived closed-form pricing expressions for VIX derivatives and computed the characteristic function through a system of ODEs. This allowed for semi-analytic pricing of both SPX and VIX options, entirely avoiding Monte Carlo simulations.
- We conducted a comparative study of the three models using real market data. A multi-step calibration cascade was implemented to sequentially align discount curves, forward prices, and options smiles. This procedure revealed both the capabilities and limitations of each model in reproducing observed SPX and VIX market characteristics in a consistent and computationally efficient manner. The results demonstrate that even structurally parsimonious models—requiring only five to seven parameters—can deliver meaningful calibration performance while maintaining tractability and interpretability.

Future directions may include extending the CAR model to higher orders (e.g., CAR(3) or CAR(p)), introducing time-dependent parameters, or exploring the estimation of autoregressive coefficients using econometric techniques on historical volatility data. Such approaches could enhance the model's empirical foundation and provide better starting points for calibration, further improving stability and accuracy.

Bibliography

Azzone, M. and Baviera, R. (2021). Synthetic forwards and cost of funding in the equity derivative market. *Finance Research Letters*, 41:101841.

Azzone, M. and Baviera, R. (2023). A fast Monte Carlo scheme for additive processes and option pricing. *Computational Management Science*, 20(1):31.

Barndorff-Nielsen, O. E. (1997). Normal Inverse Gaussian Distributions and Stochastic Volatility Modelling. *Scandinavian Journal of statistics*, 24(1):1–13.

Benth, F. E., Benth, J. S., and Koekebakker, S. (2008). *Stochastic Modelling of Electricity and Related Markets*, volume 11. World Scientific.

Bergomi, L. (2008). Smile dynamics III. *Available at SSRN 1493308*.

Brockwell, P. J. (2001a). Continuous-time ARMA processes. *Handbook of statistics*, 19:249–276.

Brockwell, P. J. (2001b). Lévy-Driven CARMA processes. *Annals of the Institute of Statistical Mathematics*, 53(1):113–124.

Brockwell, P. J. and Lindner, A. (2009). Existence and uniqueness of stationary Lévy-driven CARMA processes. *Stochastic processes and their applications*, 119(8):2660–2681.

Cont, R. and Kokholm, T. (2013). A consistent pricing model for index options and volatility derivatives. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 23(2):248–274.

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.

Cuchiero, C., Gazzani, G., Möller, J., and Svaluto-Ferro, S. (2024). Joint calibration to SPX and VIX options with signature-based models. *Mathematical Finance*, 35(1):161–213.

Duffie, D., Pan, J., and Singleton, K. (2000). Transform Analysis and Asset Pricing for Affine Jump-diffusions. *Econometrica*, 68(6):1343–1376.

Engan, A. (2013). Numerical Option Pricing in CARMA Models. Master's thesis.

Gatheral, J. (2008). Consistent modeling of SPX and VIX options. In *Bachelier congress*.

Guyon, J. (2020). The joint S&P 500/VIX smile calibration puzzle solved. *Risk, April*, 1.

Heston, S. L. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The review of Financial Studies*, 6(2):327–343.

Heston, S. L. and Nandi, S. (1998). Preference-free option pricing with path-dependent volatility: A closed-form approach. Technical report, Working Paper.

Higham, N. J. (2008). *Functions of matrices: theory and computation*. SIAM.

Jaber, E. A., Illand, C., et al. (2023). The quintic Ornstein-Uhlenbeck volatility model that jointly calibrates SPX & VIX smiles. *arXiv preprint arXiv:2212.10917*.

Kokholm, T. and Stisen, M. (2015). Joint pricing of VIX and SPX options with stochastic volatility and jump models. *The Journal of Risk Finance*, 16(1):27–48.

Lewis, A. L. (2001). A Simple Option Formula for General Jump-diffusion and Other Exponential Lévy Processes. *Available at SSRN 282110*.

Lukacs, E. (1972). A survey of the theory of characteristic functions. *Advances in Applied Probability*, 4(1):1–37.

Maffei, C. (2020). On the joint calibration of S&P 500 and VIX smiles: a challenge between two models.

Mechkov, S. (2015). Fast-reversion limit of the Heston model. *Available at SSRN 2418631*.

Schöbel, R. and Zhu, J. (1999). Stochastic Volatility With an Ornstein–Uhlenbeck Process: An Extension. *Review of Finance*, 3(1):23–46.

Song, Z. and Xiu, D. (2012). A tale of two option markets: State-price densities implied from S&P 500 and VIX option prices. *Unpublished working paper. Federal Reserve Board and University of Chicago*.

Tauchen, G. and Todorov, V. (2004). Simulation methods for Lévy-driven CARMA stochastic volatility models. *Available at SSRN 591326*.

Uhlenbeck, G. E. and Ornstein, L. S. (1930). On the theory of the Brownian motion. *Physical review*, 36(5):823.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of financial economics*, 5(2):177–188.

A | Appendix: Proofs

A.1. Proof of Heston ODE

In this Appendix we solve the function $C(t)$ in (2.6).

$$\dot{C} \;=\; C^2\left(-\frac{1}{2}\gamma^2 a\right) \;+\; C\left(-ipa\gamma\sigma\rho + a\right) \;+\; \frac{1}{2}a\sigma^2 ip(1-ip) \;\; . \tag{A.1}$$

We rewrite equation (A.1) in a general setting as

$$\dot{\alpha}(t) \;=\; A\alpha(t) \;+\; B\alpha(t)^2 \;+\; C \;\; , \tag{A.2}$$

where

$$A = -iua\gamma\sigma\rho + a \;\; ; \qquad B = -\frac{1}{2}\gamma^2 a \;\; ; \qquad C = \frac{1}{2}a\sigma^2 iu(1-iu) \;\; .$$

We integrate

$$\begin{aligned}
\int \frac{\dot{\alpha(t)}}{A\alpha(t) + B\alpha(t)^2 + C}\, dt \;&=\; \int dt = t \;+\; const \;\; . \\
\int \frac{1}{A\alpha + B\alpha^2 + C}\, d\alpha \;&=\; \int \frac{1}{B\left(\alpha + \frac{A}{2B}\right)^2 + \left(C - \frac{A^2}{4B}\right)}\, d\alpha \;=\; \left\{ \begin{array}{c} u = \alpha + \dfrac{A}{2B} \\ du = d\alpha \end{array} \right\} \\
&= \int \frac{1}{Bu^2 + (C - \frac{A^2}{4B})}\, du \;=\; \left\{ z = C - \frac{A^2}{4B} \right\} \\
&= \frac{1}{B}\int \frac{1}{u^2 + z/B}\, du \\
&= \frac{1}{B\sqrt{z/B}} \arctan\left(\frac{u}{\sqrt{z/B}}\right) + const \\
&= \frac{1}{\sqrt{CB - \frac{A^2}{4}}} \arctan\left\{ \frac{\alpha + \frac{A}{2B}}{\sqrt{\frac{C}{B} - \frac{A^2}{4B^2}}} \right\} + const \;\; .
\end{aligned}$$

$$\Rightarrow \arctan\left\{\frac{\alpha+\frac{A}{2B}}{\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}}}\right\} = \sqrt{CB-\frac{A^2}{4B}}\,(t+const) \quad .$$

$$\Rightarrow \alpha(t) = -\frac{A}{2B}+\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}}\tan\left\{B\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}}\,t+const\right\} \quad ,$$

with terminal condition $\alpha(T)=0$.

$$const = \arctan\left\{\frac{\frac{A}{2B}}{\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}}}\right\} - B\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}}\,T \quad .$$

$$\Rightarrow \alpha(t) = -\frac{A}{2B}+\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}}\tan\left\{B\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}}\,(t-T)+\arctan\left[\frac{\frac{A}{2B}}{\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}}}\right]\right\} \quad ,$$

which we can rewrite as

$$\alpha(t) = D+E\tan\left(\arctan\left(-\frac{D}{E}\right)-F(T-t)\right) \quad , \tag{A.3}$$

with

$$D=-\frac{A}{2B} \;\; ; \qquad E=\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}} \;\; ; \qquad F=B\sqrt{\frac{C}{B}-\frac{A^2}{4B^2}} \quad .$$

We simplify the expression with

$$
\begin{aligned}
1)\quad & \cos\left(\arctan\left(-\frac{D}{E}\right) - F(T-t)\right) \\
&= \cos\left(\arctan\left(-\frac{D}{E}\right)\right)\cos\left(F(T-t)\right) \;+\; \sin\left(\arctan\left(-\frac{D}{E}\right)\right)\sin\left(F(T-t)\right) \\
&= \frac{E}{\sqrt{D^2+E^2}}\cos\left(F(T-t)\right) \;-\; \frac{D}{\sqrt{D^2+E^2}}\sin\left(F(T-t)\right) \\
&= \frac{E\cos\left(F(T-t)\right) \;-\; D\sin\left(F(T-t)\right)}{\sqrt{D^2+E^2}} \;\;;
\end{aligned}
$$

$$
\begin{aligned}
2)\quad & \sin\left(\arctan\left(-\frac{D}{E}\right) - F(T-t)\right) \\
&= \sin\left(\arctan\left(-\frac{D}{E}\right)\right)\cos\left(F(T-t)\right) \;-\; \cos\left(\arctan\left(-\frac{D}{E}\right)\right)\sin\left(F(T-t)\right) \\
&= -\frac{D\cos\left(F(T-t)\right) \;+\; E\sin\left(F(T-t)\right)}{\sqrt{D^2+E^2}} \;\;;
\end{aligned}
$$

$$
\begin{aligned}
3)\quad & \tan\left(\arctan\left(-\frac{D}{E}\right) - F(T-t)\right) \\
&= \frac{\sin(\dots)}{\cos(\dots)} \;=\; -\frac{D\cos\left(F(T-t)\right) \;+\; E\sin\left(F(T-t)\right)}{E\cos\left(F(T-t)\right) \;-\; D\sin\left(F(T-t)\right)} \\
&= -\frac{D\cot\left(F(T-t)\right) \;+\; E}{E\cot\left(F(T-t)\right) \;-\; D} \;\;.
\end{aligned}
$$

Yielding the expression

$$
\Rightarrow \alpha(t) \;=\; D \;+\; E\left[\frac{-D\cot\left(F(T-t)\right) - E}{E\cot\left(F(T-t)\right) - D}\right] \;=\; \frac{-\left(D^2+E^2\right)}{E\cot\left(F(T-t)\right) - D} \;\;. \tag{A.4}
$$

Plugging the original parameters in the solution in (A.4), for $\alpha(t) = C(t)$

$$
\begin{aligned}
D \;&=\; \frac{iua\gamma\sigma\rho - a}{-\gamma^2 a} \;=\; \frac{1 - iu\sigma\rho\gamma}{\gamma^2} \;\;, \\
E \;&=\; \sqrt{\frac{\frac{1}{2}a\sigma^2 iu(1-iu)}{\frac{1}{2}\gamma^2 a} \;-\; \frac{(a - iua\gamma\sigma\rho)^2}{\gamma^4 a^2}} \;=\; \frac{i}{\gamma}\sqrt{\sigma^2 iu(1-iu) + (\gamma^{-1} - \rho\sigma iu)^2} \;=\; \frac{i}{\gamma}\Omega(u) \;\;, \\
F \;&=\; -\frac{1}{2}\gamma^2 aE \;=\; -\frac{i}{2}\gamma a\Omega(u) \;\;,
\end{aligned}
$$

with the notation

$$\Omega(u) \;=\; \sqrt{\sigma^2 iu(1-iu) + (\gamma^{-1} - \rho\sigma iu)^2} \;, \qquad \Psi(u) \;=\; \frac{\gamma^{-1} - iu\rho\sigma}{\Omega(u)} \;. \tag{A.5}$$

$$D^2 + E^2 = \frac{\cancel{(iu\sigma\rho - \gamma^{-1})^2}}{\gamma^2} + \left\{-\gamma^{-2}\left[\sigma^2 iu(1-iu) + \cancel{(iu\sigma\rho - \gamma^{-1})^2}\right]\right\} = -\frac{\sigma^2 iu(1-iu)}{\gamma^2} \;.$$

$$\begin{aligned}
\Rightarrow C(t) &= \frac{\frac{\sigma^2 iu(1-iu)}{\gamma^2}}{i\frac{\Omega(u)}{\gamma}\cot\left\{-\frac{i}{2}\gamma a\Omega(u)(T-t)\right\} - \frac{1-iu\sigma\rho\gamma}{\gamma^2}} \\
&= \frac{\sigma^2 iu(1-iu)}{i\gamma\Omega(u)\cot\left\{-\frac{i}{2}\gamma a\Omega(u)(T-t)\right\} - 1 + iu\sigma\rho\gamma} \\
&= \frac{\sigma^2 iu(1-iu)}{i^2\gamma\Omega(u)\coth\left\{\frac{1}{2}\gamma a\Omega(u)(T-t)\right\} - 1 + iu\sigma\rho\gamma} \\
&= -\frac{\gamma^{-1}\Omega(u)(1-\Psi(u)^2)}{\coth\left\{\frac{1}{2}a\gamma\Omega(u)(T-t)\right\} + \Psi(u)} \;.
\end{aligned}$$

A.2. Proof of Heston Characteristic Function

In this Appendix we prove the solution (2.8).
We take the solution of $C(t)$ in the form (A.3)

$$\begin{aligned}
\dot{A}(t) \;&=\; -C(t) \;=\; -D - E\tan\left\{\arctan\left(-\frac{D}{E}\right) - F(T-t)\right\} \;. \\
A(t) \;&=\; A(0) - Dt - E\int_0^t \tan\left\{\arctan\left(-\frac{D}{E}\right) - F(T-t)\right\} ds \\
&=\; A(0) - Dt - \frac{E}{F}\log\left\{\frac{\cos\left(\arctan\left(-\frac{D}{E}\right) - FT\right)}{\cos\left(\arctan\left(-\frac{D}{E}\right) - F(T-t)\right)}\right\} \;.
\end{aligned}$$

Terminal condition $A(T) = 0$ implies that

$$A(O) \;=\; DT + \frac{E}{F}\log\left\{\frac{\sqrt{E^2+D^2}}{E}\cos\left(\arctan\left(-\frac{D}{E}\right) - FT\right)\right\} \;. \tag{A.6}$$

$$\begin{aligned}\Rightarrow A(t) &= D(T-t) + \frac{E}{F}\log\left\{\frac{\sqrt{E^2+D^2}}{E}\cos\left(\arctan\left(-\frac{D}{E}\right) - F(T-t)\right)\right\} \\ &= D(T-t) + \frac{E}{F}\log\left\{\cos\left[F(T-t)\right] - \frac{D}{E}\sin\left[F(T-t)\right]\right\} \quad .\end{aligned} \tag{A.7}$$

Plugging the original parameter in (A.7), and following the notation in (A.5):

$$D = \frac{1-iu\sigma\rho\gamma}{\gamma^2} \quad , \qquad E = i\Omega(u)\gamma^{-1} \quad , \qquad F = -\frac{i}{2}a\gamma\Omega(u) \quad , \tag{A.8}$$

we finally get

$$\begin{aligned}A(t) &= \frac{1-iu\sigma\rho\gamma}{\gamma^2}(T-t) \\ &- \frac{2}{a\gamma^2}\log\left\{\cos\left[-\frac{ia\gamma\Omega(u)}{2}(T-t)\right] - \frac{1-iu\sigma\rho\gamma}{i\gamma\Omega(u)}\sin\left[-\frac{ia\gamma\Omega(u)}{2}(T-t)\right]\right\} \\ &= \frac{1-iu\sigma\rho\gamma}{\gamma^2}(T-t) \\ &- \frac{2}{a\gamma^2}\log\left\{\cosh\left[\frac{a\gamma\Omega(u)}{2}(T-t)\right] + \frac{1-iu\sigma\rho\gamma}{\gamma\Omega(u)}\sinh\left[\frac{a\gamma\Omega(u)}{2}(T-t)\right]\right\} \quad ,\end{aligned}$$

and, passing to the exponential,

$$e^{A(t)} = e^{\left\{\frac{1-iu\sigma\rho\gamma}{\gamma^2}(T-t)\right\}}\left\{\cosh\left[\frac{a\gamma\Omega(u)}{2}(T-t)\right] + \frac{1-iu\sigma\rho\gamma}{\gamma\Omega(u)}\sinh\left[\frac{a\gamma\Omega(u)}{2}(T-t)\right]\right\}^{-\frac{2}{a\gamma^2}} \quad .$$

Finally, plugging $A(t)$, $B(t)$ and $C(t)$ in (2.5), we obtain the solution in (2.8)

$$\phi_T(u) = \frac{\exp\left\{\frac{-iu\sigma\rho\gamma+1}{\gamma^2}T\right\}\exp\left\{\frac{-\gamma^{-1}\sigma^2 iu(1-iu)\Omega^{-1}}{\coth\left(\frac{a\gamma\sigma}{2}T\right)+\Psi}(a^{-1}z_0)\right\}}{\left[\cosh\left(\frac{a\gamma\sigma}{2}T\right) + \frac{1-iu\sigma\rho\gamma}{\gamma\Omega}\sinh\left(\frac{a\gamma\sigma}{2}T\right)\right]^{\frac{2}{a\gamma^2}}} \quad . \tag{A.9}$$

B | Appendix: Plots

In this Appendix, we present the calibration results for the July 20, 2021 dataset.

B.1. Heston

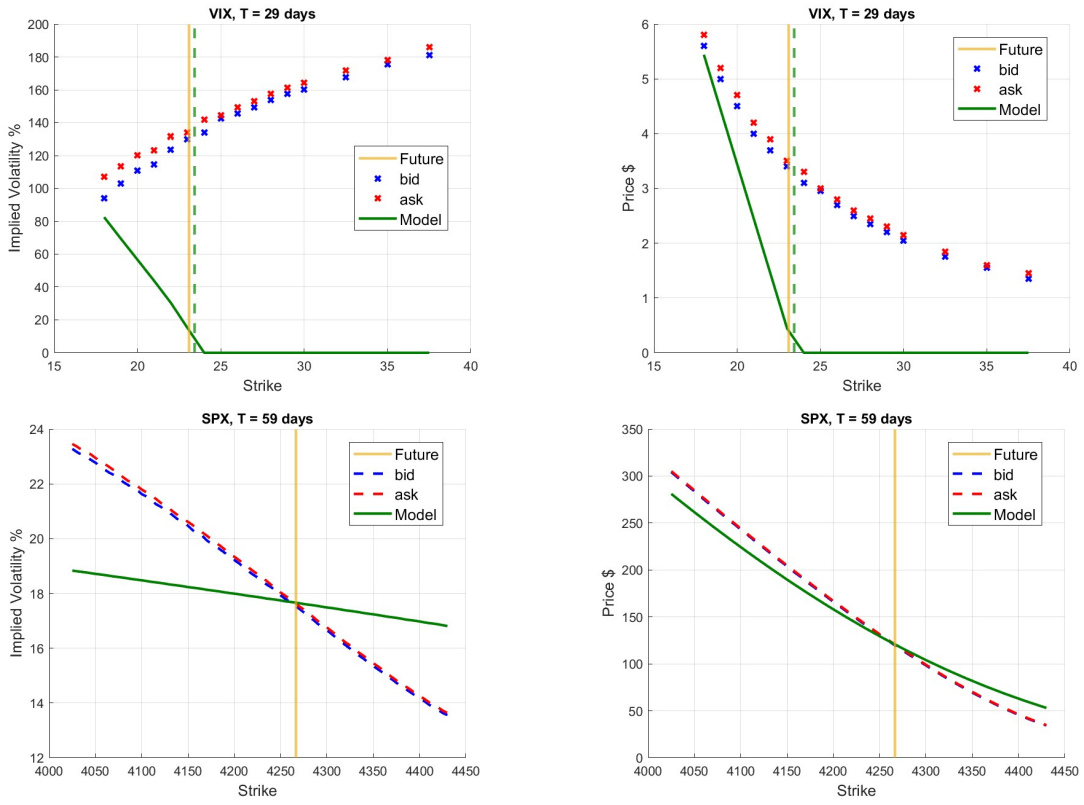

Figure B.1: Maturity slice T_1: SPX-VIX joint calibration results of Heston model on July 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

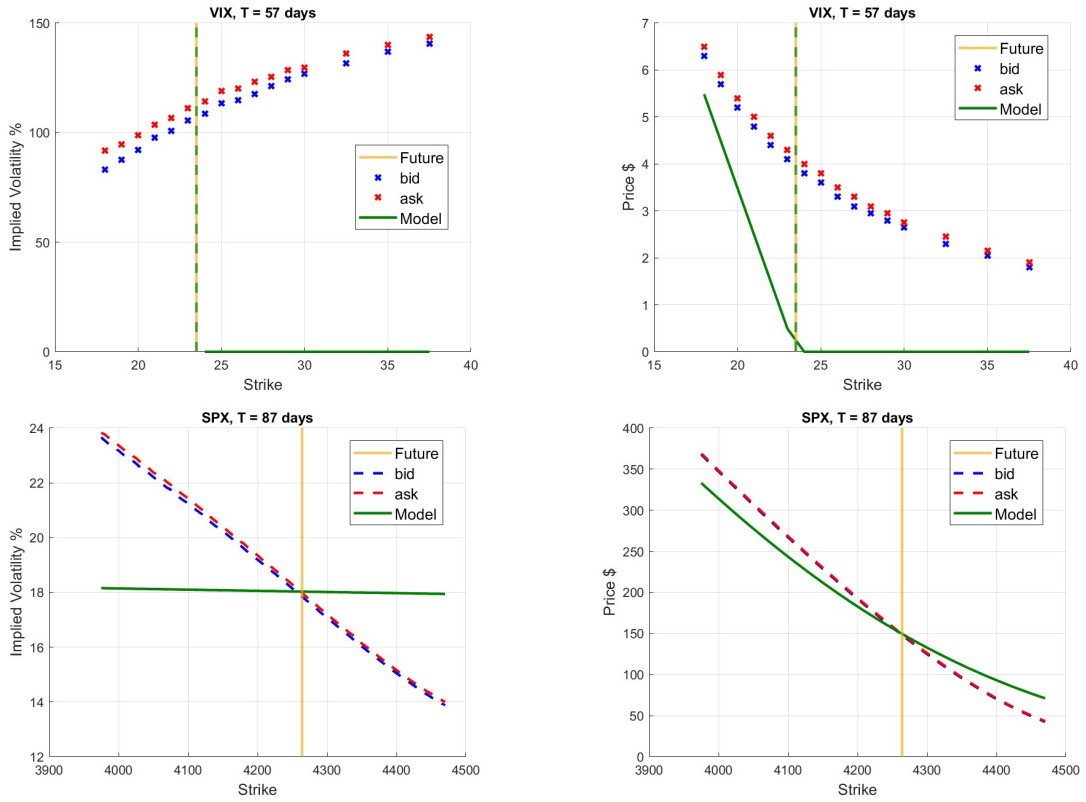

Figure B.2: Maturity slice T_2: SPX-VIX joint calibration results of Heston model on July 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

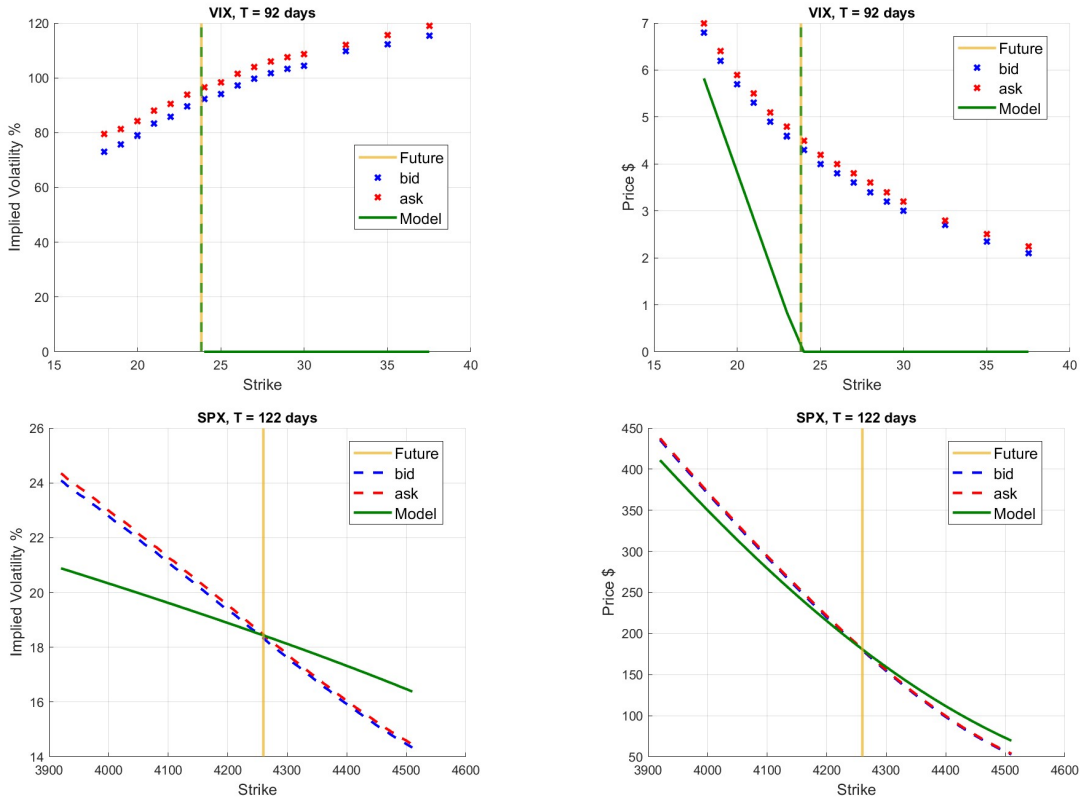

Figure B.3: Maturity slice T_3: SPX-VIX joint calibration results of Heston model on July 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

Heston model calibrated parameters

	σ	γ	ρ	a	x_0
July 20, 2021 dataset					
$T_{\text{VIX}} = 29\text{d}$	0.7288	1.7908	-1.0000	0.1238	0.0501
$T_{\text{VIX}} = 57\text{d}$	0.7629	1.7902	-1.0000	0.0096	0.0548
$T_{\text{VIX}} = 92\text{d}$	0.6490	1.7903	-1.0000	0.2555	0.0467

Table B.1: Heston model VIX-SPX jointly calibrated parameters for the July 2021 dataset.

B.2. Ornstein-Uhlenbeck

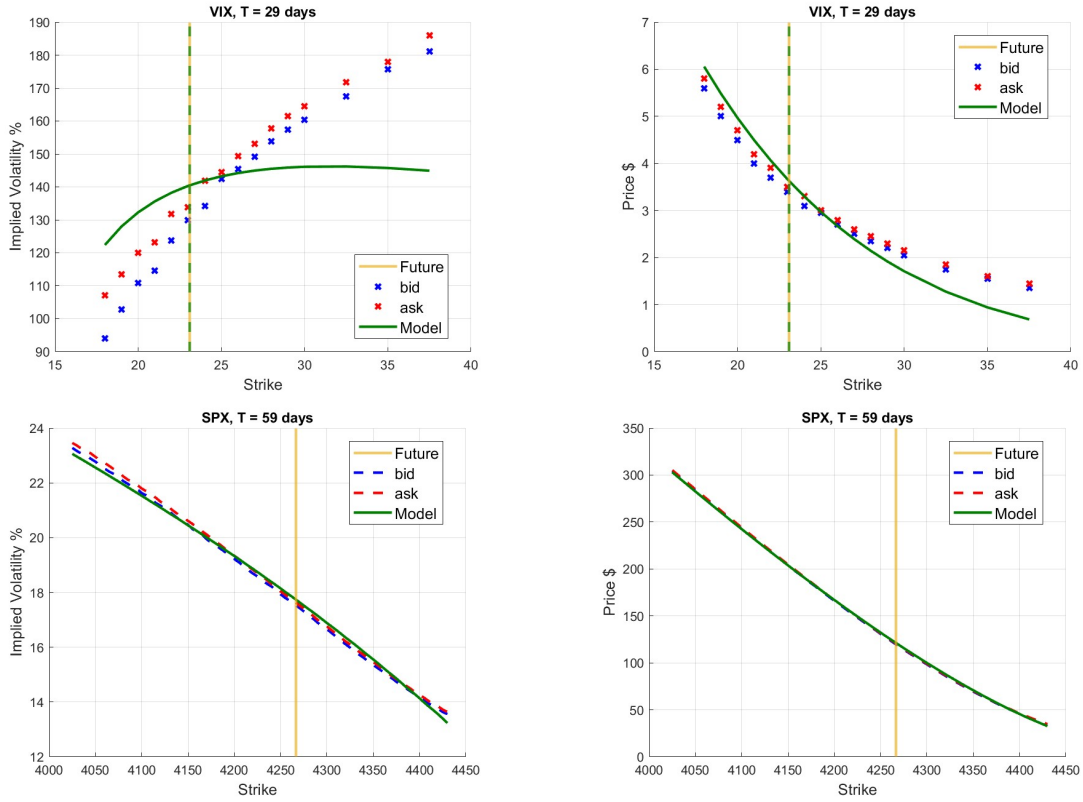

Figure B.4: Maturity slice T_1: SPX-VIX joint calibration results of OU model on July 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

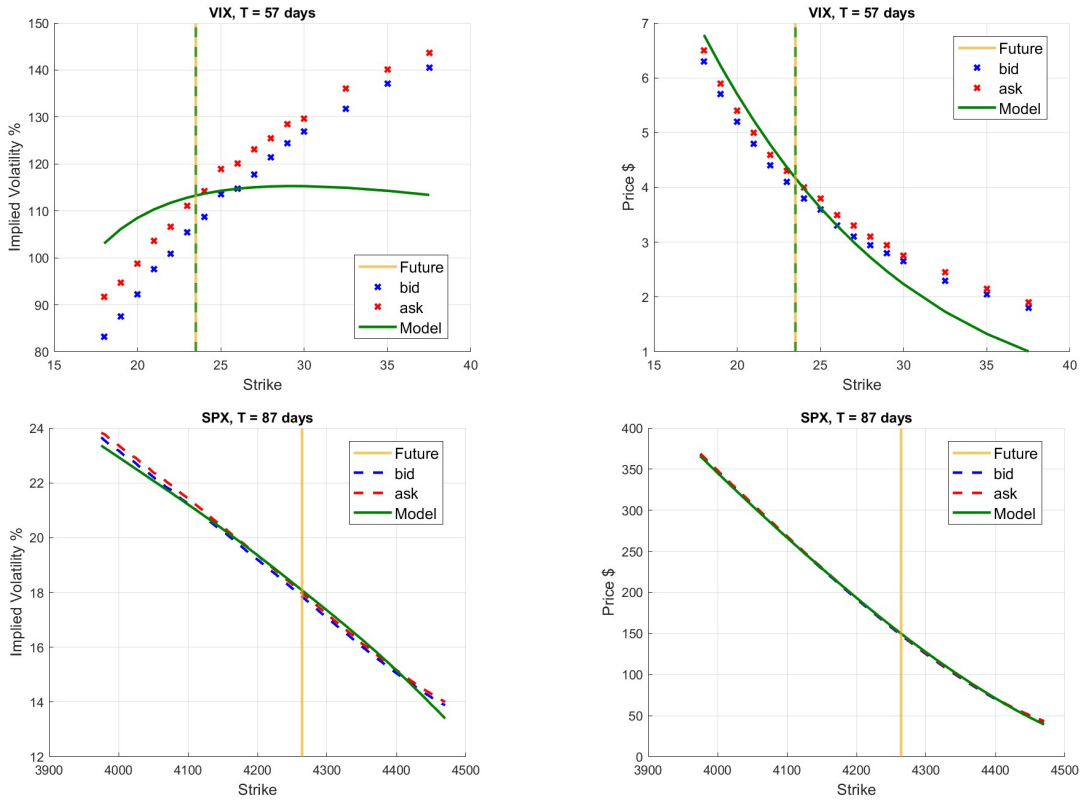

Figure B.5: Maturity slice T_2: SPX-VIX joint calibration results of OU model on July 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

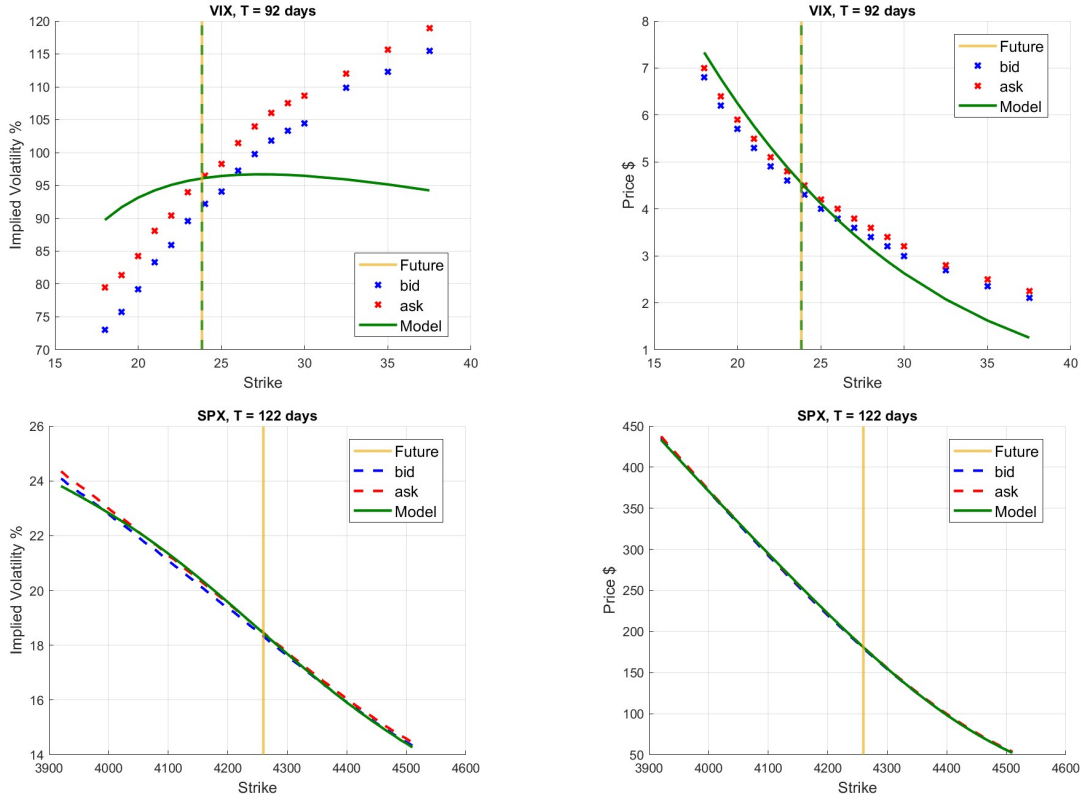

Figure B.6: Maturity slice T_3: SPX-VIX joint calibration results of OU model on July 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

OU model calibrated parameters

	μ	σ	γ	ρ	a	x_0
July 20, 2021 dataset						
$T_{\text{VIX}} = 29\text{d}$	0.1714	1.0060	0.7300	-0.9428	1.0601	0.0617
$T_{\text{VIX}} = 57\text{d}$	0.0804	1.0027	0.6993	-0.9654	2.4687	0.0670
$T_{\text{VIX}} = 92\text{d}$	0.0713	0.9581	0.6646	-0.9716	2.9250	0.1189

Table B.2: OU model VIX-SPX jointly calibrated parameters for the July 2021 dataset.

B.3. Continuous Autoregressive

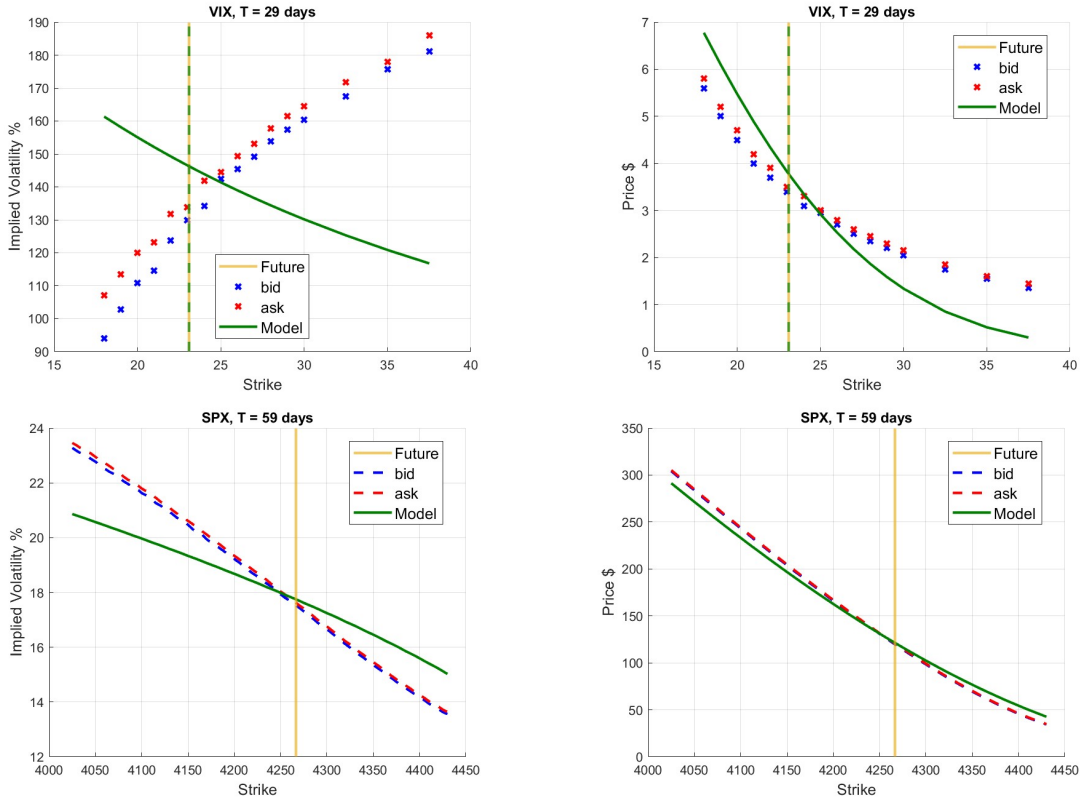

Figure B.7: Maturity slice T_1: SPX-VIX joint calibration results of CAR(2) model on July 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

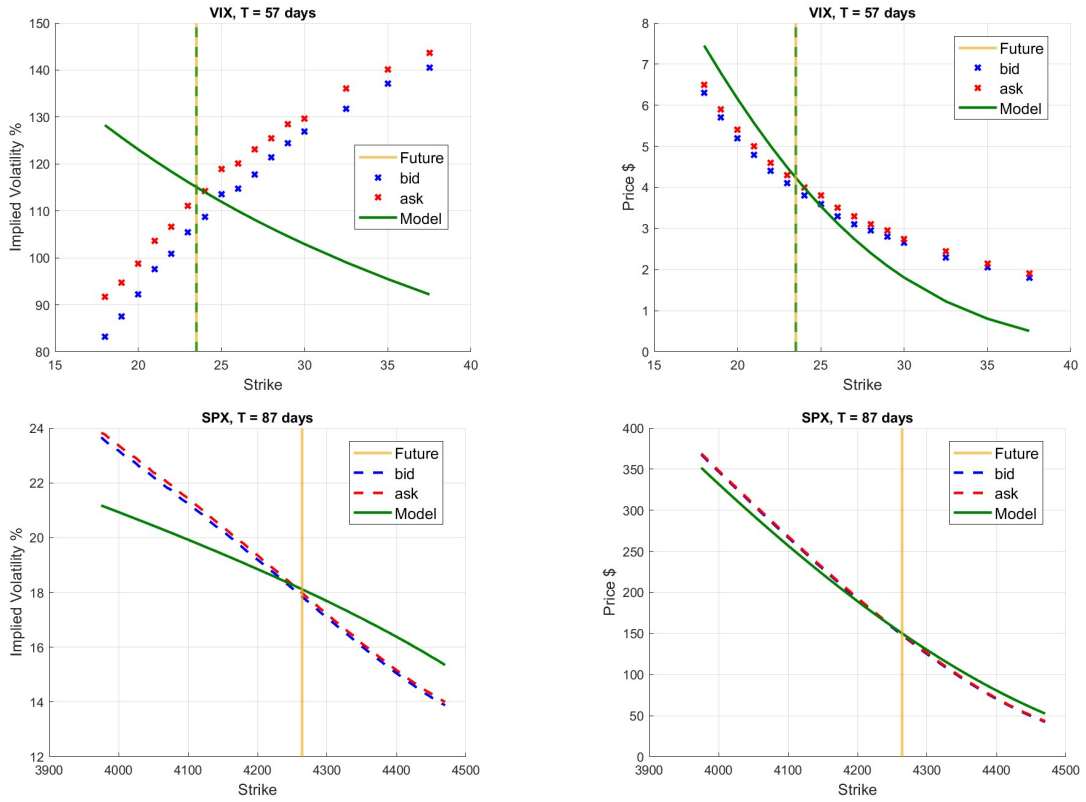

Figure B.8: Maturity slice T_2: SPX-VIX joint calibration results of CAR(2) model on July 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

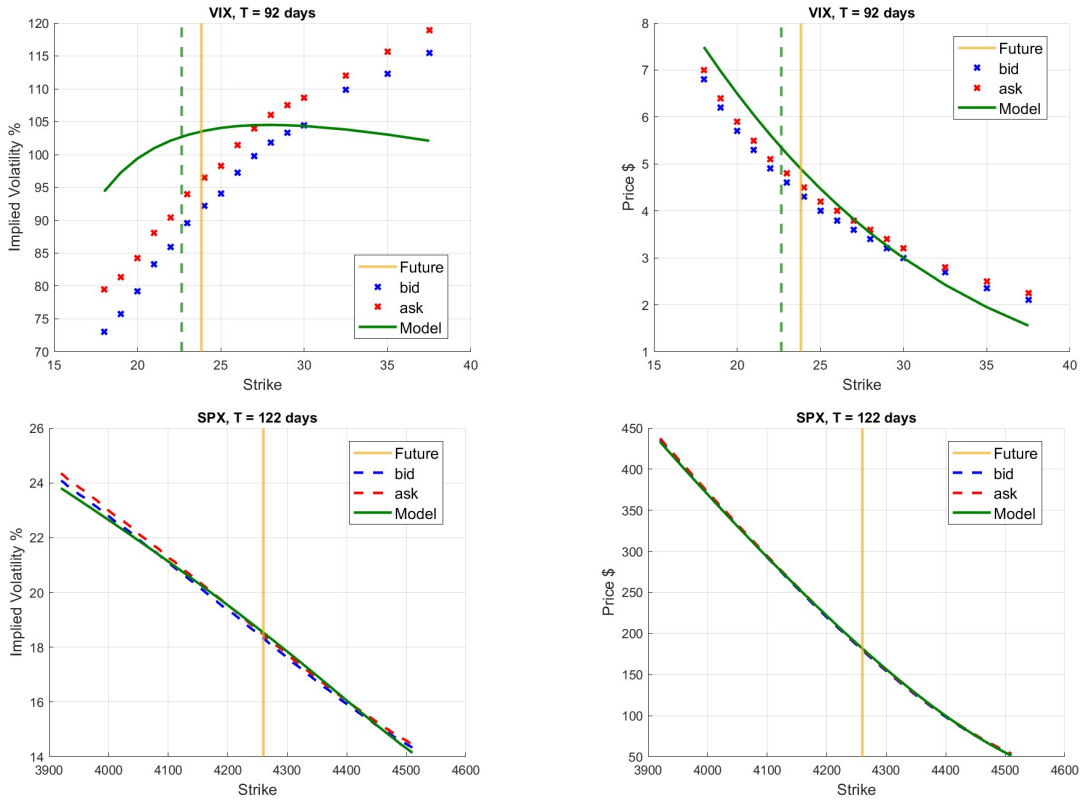

Figure B.9: Maturity slice T_3: SPX-VIX joint calibration results of CAR(2) model on July 20, 2021 data. Top panels display VIX option market quotes—implied volatilities (left) and option prices (right)—as functions of strike, with bid/ask levels in blue/red and model output in green. Bottom panels show the same for SPX options. The vertical gold line marks the observed future levels, while the vertical green dashed line indicates the model-implied VIX future level.

CAR(2) model calibrated parameters

	σ	γ	ρ	α_1	α_2	x_1	x_2
July 20, 2021 dataset							
$T_{\text{VIX}} = 29\text{d}$	0.7839	20.0218	-1.0000	52.2996	16.22446	0.1896	2.0043
$T_{\text{VIX}} = 57\text{d}$	0.7852	14.7361	-1.0000	49.9284	8.5201	0.1569	3.8720
$T_{\text{VIX}} = 92\text{d}$	1.1137	34.5114	-0.9818	59.9723	59.9619	0.0991	-5.8981

Table B.3: CAR(2) model VIX-SPX jointly calibrated parameters for the July 2021 dataset.

C | Appendix: Codes

C.1. Runge-Kutta scheme

```
function [out, alpha, beta] = cfCAR(u, x0, tau, gamma, sigma,
    rho, a1, a2, nsteps)
%CFCAR Computes the characteristic function of a CAR2 process
   .
%
%   W = CFCAR(U, X0, TAU, GAMMA, SIGMA, RHO, A1, A2)
%   Computes the characteristic function for a purely
   diffusive affine process.
%
%   Inputs:
%   - U: Frequency variable (Nx1 or 1xN)
%   - X0: Initial state (Nx1 vector)
%   - TAU: Time horizon (scalar)
%   - GAMMA, SIGMA, RHO: CAR parameters (scalars)
%   - A1, A2: Alpha 1 and 2 in the matrix A (negative, real
   scalars)
%
%   Outputs:
%   - OUT: Characteristic function value at U
%   - ALPHA, BETA: Riccati equation solutions

% Ensure u is a row vector (1xNu)
u = reshape(u, 1, []) * 1i;
nu = length(u);

% Initialize beta (NxNu) and alpha (1xNu)
nx = length(x0);
```

```
beta = [u; zeros(nx - 1, nu)];
alpha = zeros(1, nu);

% Time step for numerical integration
if nargin < 9
    nyear = 200;  % or set t to a specific default value
else
    nyear = nsteps;
end

dt = tau / (ceil(tau * nyear));

% Define Riccati system
riccati = @(beta) cf_riccati(beta, gamma, sigma, rho, a1, a2,
    nx);

% Runge-Kutta 4th order integration
for k = tau - dt : -dt : 0
    K1 = riccati(beta);
    beta1 = beta + (dt / 2) * K1(2:end, :);
    K2 = riccati(beta1);
    beta2 = beta + (dt / 2) * K2(2:end, :);
    K3 = riccati(beta2);
    beta3 = beta + dt * K3(2:end, :);
    K4 = riccati(beta3);

    alpha = alpha + (dt / 6) * (K1(1, :) + 2 * K2(1, :) + 2 *
        K3(1, :) + K4(1, :));
    beta = beta + (dt / 6) * (K1(2:end, :) + 2 * K2(2:end, :)
        + 2 * K3(2:end, :) + K4(2:end, :));
end

% Compute characteristic function
out = exp(alpha.' + beta.' * x0);
% Preserve the original shape of u
out = reshape(out, size(u));
```

```
end

% Riccati System
function out = cf_riccati(beta, gamma, sigma, rho, a1, a2, nx
    )
% Solves Riccati equations
nB = size(beta, 1); % Should be nx
if nB ~= nx
    error('Dimension mismatch: beta should have %d rows.', nx
        );
end
% Compute drift contribution (Alpha)
% d /dt
out(1, :) = +gamma^2*beta(6,:) + 0.5*gamma^2*beta(3,:);
% compute Betas
% d _z/dt
out(2, :) = 0;
% d _x1/dt
out(3, :) = 0;
% d _x2/dt
out(4, :) = 0;
% d _v/dt
out(5, :) = -(0.5*sigma^2*beta(1,:) + a2*beta(5,:) - 0.5*
    sigma^2*beta(1,:).^2 - rho*sigma*gamma*beta(1,:).*beta
    (5,:) - 0.5*gamma^2*beta(5,:).^2);
% d _w/dt
out(6, :) = -(-2*beta(4,:) + a1*beta(5,:) + 2*a2*beta(6,:) -
    2*rho*sigma*gamma*beta(1,:).*beta(6,:) - 2*gamma^2*beta
    (5,:).*beta(6,:));
% d _r/dt
out(7, :) = -(-beta(5,:) + 2*a1*beta(6,:) - 2*gamma^2*beta
    (6,:).^2);
end
```

List of Figures

List of Tables

List of Symbols

Symbol	Description
$\mathbb{E}$	Expected value
$\mathbb{R}$	Real numbers
$\mathbb{R}^+$	Non-negative real numbers
$\mathbb{C}$	Complex numbers
$\mathcal{N}(\mu, \sigma^2)$	Gaussian distribution with mean μ and variance σ^2
$\Gamma(\alpha, \theta)$	Gamma distribution with shape α and scale θ
$\chi^2_\nu(\delta)$	Noncentral Chi-squared with degrees of freedom ν and noncentrality parameter δ
λ_i	Eigenvalues
$\mathcal{F}_T$	Filtration at time T
$\xi_T(u)$	Forward variance process curve at time T for maturity u, $T \leq u$
Δ	30 days time interval for VIX calculation
$V(\cdot)$	Derivative payoff
$\phi_{s,t}$	Characteristic Function of increment between s and t time to maturity
ϕ_t	Characteristic Function at time to maturity t
$F(t_0, t)$	Forward price with maturity t at time t_0
$B(t_0, t)$	Discount factor between t_0 and t
θ	Set of model parameters
$J(\theta)$	Loss function evaluated using parameter set θ
$W(t)$	Standard Brownian motion at time t
μ	Long-term mean of volatility process
a	Speed of mean-reversion
ρ	Correlation coefficient between Brownian motions
σ	Volatility parameter
γ	Volatility-of-volatility parameter
α_i	Matrix A parameters in CAR(2) model with $i = 1, 2$

List of Shorthands

Symbol	Description
VIX	Cboe Volatility Index
SPX	S&P 500 Index
FRH	Fast Reversion Heston
NIG	Normal Inverse Gaussian
OU	Ornstein–Uhlenbeck
CARMA	Continuous Autoregressive Moving Average
CAR	Continuous Autoregressive
AR	Autoregressive
ODE	Ordinary Differential Equation
PDE	Partial Differential Equation
SDE	Stochastic Differential Equation
FFT	Fast Fourier transform

Acknowledgements

I am deeply grateful to my advisors, Prof. Baviera and Prof. Azzone, for their invaluable guidance, expertise, and support throughout this journey.

I would like to sincerely thank my mum, dad, brother, and the rest of my family for their constant encouragement over the past years.

A heartfelt thank you to Maria Elena—for her patience, understanding, and for always being there for me and standing by my side.

Lastly, a big thank you to all my friends, who have been a constant source of motivation, support, and laughter.