RISK MANAGEMENT MAGAZINE

Vol. 18, Issue 2
May – August 2023

EXCERPT

https://www.aifirm.it/rivista/progetto-editoriale/

Analysis of numerical integration schemes for the Heston model: a case study based on the pricing of investment certificates

Michelangelo Fusaro; Pier Giuseppe Giribone; Alessio Tissone

Analysis of numerical integration schemes for the Heston model: a case study based on the pricing of investment certificates

Michelangelo Fusaro (Quantitative Financial Analyst and AIAF/AIFIRM member) – Pier Giuseppe Giribone (BPER Banca – Financial Engineering and University of Genoa – Department of Economics) Alessio Tissone (Quantitative Financial Analyst and AIAF/AIFIRM member)

Corresponding Author: Pier Giuseppe Giribone (piergiuseppe.giribone@bper.it)

Article submitted to double-blind peer review, received on 29th July 2023 and accepted on 7th August 2023

Abstract

The Heston model is one of the most used techniques for estimating the fair value and the risk measures associated with investment certificates. Typically, the pricing engine implements a significant number of projections of the underlying until maturity, it calculates the pay-off for all the paths thus simulated considering the characteristics of the structured product and, in accordance with the Monte Carlo methodology, it determines its theoretical value by calculating its mean and discounting it at valuation time. In order to generate the future paths, the two stochastic differential equations governing the dynamics of the Heston model should be integrated simultaneously over time: both the one directly associated with the underlying and the one associated with variance. Consequently, it is essential to implement a numerical integration scheme that allows such prospective simulations to be implemented. The present study aims to consider alternatives to the traditional Euler method with the aim of reducing or in some cases eliminating the probability of incurring unfeasible simulated values for the variance. In fact, one of the main drawbacks of the Euler basic integration scheme applied to the Heston bivariate stochastic model is that of potentially generating negative variances in the simulation that should be programmatically corrected each time such undesired effect occurs. The methods which do not intrinsically admit the generation of negative values of the variance proved to be very interesting, in particular the Transformed Volatility scheme.

Key Words:
Certificate pricing, Stochastic Differential Equation, Heston model, Numerical integration schemes.

JEL code: G12-G17-C53-C63

1) Introduction

The financial industry has paid particular attention to the issue of structured products and investment certificates in order to meet the needs of investors, as shown in the analysis of the primary market made by ACEPI (Italian Association of Investment Certificates and Products). The first quarter of 2023 showed that the total volumes, equal to 5,572 million euros issued by ACEPI members, increased by 37% compared to the previous quarter and compared to the average of the quarterly issues in 2022. The increase in amounts issued in Q1 2023 strengthens the growth trend recorded in the second half of 2022 (the year ended overall at a figure of 16,236 million euros), up 71% compared to 2021. In a context of uncertainty regarding monetary policies, market trends and the stability of the financial system, the search for protection was one of the factors explaining the growing trend. The number of products offered, 376, exceeded the record of 343 already achieved in Q4 last year, confirming a further 10% increase. In terms of breakdown into ACEPI macro-classes, in Q1 protected-capital products accounted for 63% of issues on the primary market compared to 30% of conditionally protected capital products. The remaining 7% refers to Credit Linked Notes, a type of product which had shown considerable growth in 2022, with peaks of 11% and 10% of the total issues in Q3 and Q4.

The figures presented briefly in this context suggest the importance for financial institutions to estimate the fair value and the sensitivity measures of these products as precisely and reliably as possible. To this end, after focusing on variance reduction techniques (Bottasso *et al.*, 2023a) and on low discrepancy sequences (Bottasso *et al.*, 2023b), the present study centers on the numerical integration schemes associated with one of the most widespread quantitative analysis models for investment certificates: the Heston model (1993). In particular, the goal is to look for valid alternatives to the most widespread method of numerical integration, i.e., the Euler method which, due to the type of stochastic process, could unfortunately present the drawback of making the variance assume negative and, consequently, not eligible values (Rouah, 2013). Such variance values, unacceptable from a theoretical point of view, are artificially set to zero or considered for their absolute value. Thus, the importance of studying more robust integration schemes that minimize (or, in some cases, eliminate) the probability that this undesired effect occurs. The problem is well known in the scientific literature, in fact several studies applied to different types of options have been dedicated to this research, including: Vanilla options (Mrázek and Pospíšil, 2017), Asian options (Begin, Bedard and Gaillardetz, 2015), Forward-Start options (Broadie and Kaya, 2006), Double-no-touch options (Lord, Koekkoek and van Dijk, 2008). On the other hand, the implementation of other schemes, different from the traditional Euler scheme, in the pricing of hybrid instruments, i.e., characterized by a Fixed Income component and one or more option strategies, such as certificates, is innovative.

In the next section of the paper we will describe the integration schemes that allow to reduce the numerical approximation error of the dynamics. They will then be implemented in the third section of the paper applied to the pricing of a plain vanilla option written on the FTSE MIB index using market parameters. For such pricing problem, an analytical valuation formula is known, i.e. the price of a European call option can be expressed in a closed form involving integrals in the complex numbers that can be numerically valuated with the Gauss-Laguerre quadrature. Having a reasonably exact expression to which the Monte Carlo method should converge makes it possible to test the stability (absence of convergence bias to the fair value) and measure the performance across numerical integration schemes. The schemes that proved to be most stable and efficient have been taken into consideration to be implemented in a valuation context for which an analytical pricing formula is not available. The last section of the study, therefore, implements the schemes that resulted most suitable in the previous tests and values, again using market parameters, the most widespread types of investment certificates, in accordance with the ACEPI statistics previously presented: i.e. products characterized by digital coupons with or without memory, autocallability and conditionally protected capital.

2) An overview of the numerical integration schemes for the Heston model

The Monte Carlo method within the Heston pricing framework is understood as a set of techniques which allow the generation of an artificial historical series of prices of an underlying (typically equity or index) and of variance over time, from which option prices can be calculated. In the literature, several numerical schemes allow to achieve this goal. The first approach is to implement standard methods valid for any kind of stochastic differential equation that is supposed to be integrated over time: among these, the most popular are the Euler and the Milstein methods. The advantage of employing these multi-purpose techniques is that they are easy to understand, and their convergence properties are well known. As a result, these two schemes can be adopted for pricing a large number of financial derivatives typologies which fair values have to be estimated using a Monte Carlo technique. A second approach is to use methods designed ad-hoc for the specific Heston dynamics. As regards the Heston model, we can mention the IJK scheme of Kahl and Jäckel (2006), the transformed volatility scheme of Zhu (2010) or the moment-matching scheme of Andersen and Brotherton-Ratcliffe (2005). Such approaches, specifically applicable for the most widely used valuation model for pricing investment certificates with an equity or index underlying, are potentially able to reach a theoretically higher speed of convergence and in certain cases to avoid the unwelcome effect of generating negative variances, which might occur using a multi-purpose scheme. For a comprehensive review of methods for numerical integration, the contribution of Van Haastrecht and Pelsser (2010) should be considered.

The starting point for the study and the consequent implementation of any numerical integration technique is to consider the object of integration in a continuous form. Bearing in mind that the stock (or index) price and its variance in the Heston model are driven by the following bivariate system of stochastic differential equations (SDE):

$$\begin{aligned} dS_t &= (r-q)S_t dt + \sqrt{v_t} S_t dW_{1,t} \\ dv_t &= \kappa(\theta - v_t)dt + \sigma\sqrt{v_t} dW_{2,t} \end{aligned} \qquad (1)$$

Where $E[dW_{1,t}, dW_{2,t}] = \rho dt$

The parameters of the model are:

$(r-q)$ the drift of the process for the share or index. In particular, r is the risk-free rate and q is the dividend yield associated with the underlying. Depending on the available market data, both quantities should be time varying.

$\kappa > 0$ the mean reversion speed for the variance.

$\theta > 0$ the mean reversion level for the variance.

$\sigma > 0$ the volatility of the variance.

$v_0 > 0$ the initial level of the variance (at time zero).

Furthermore $W_{1,t}$ and $W_{2,t}$ are Wiener processes with correlation $\rho \in [-1,1]$ and S_t is the value of the price of the share/index assumed at time t. Thus, S_0 is the initial spot value.

The processes described in Equation (1) are defined in continuous time. The numerical simulation, however, must necessarily be programmed using discrete time steps. Therefore, the first step to be taken in a numerical simulation scheme is generally that of approximating the continuous-time process with a discrete-time process: in other words, discretizing the stochastic differential equations. Both the dynamics associated with the stock price and that associated with its volatility can be rewritten in the following general form, considering a generic random variable X_t:

$$dX_t = \mu(X_t, t)dt + \sigma(X_t, t)dW_t \quad (2)$$

With $\mu(X_t, t)$ being the drift and $\sigma(X_t, t)$ being the volatility of the stochastic process to be represented. X_t is simulated, along the time interval $[0, T]$, which is supposed to be divided into N points thus creating a time grid of the type $0 = t_1 < t_2 < \cdots < t_N = T$ where the time increments are characterized by the same amplitude dt. This choice also allows the use of a more convenient notation since it allows to write $t_i - t_{i-1}$ merely with dt for any $i = 2, \ldots, N$. It being understood that all results derived with equally spaced time increments can easily be extended by ranges of non-uniform amplitude. Integrating dX_t from t to $t + dt$, we have:

$$X_{t+dt} = X_t + \int_t^{t+dt} \mu(X_u, u)du + \int_t^{t+dt} \sigma(X_u, u)dW_u \quad (3)$$

Equation (3) is the starting point for the discretization. The concept is that at time t the value of X_t is known and we want to obtain the next value of the time series X_{t+dt} at time $t + dt$. Clearly, to obtain the value of an option by means of numerical simulation, which we supposed in the first part of our discussion to be a plain vanilla European option, we have to simulate the Heston bivariate process (S_t, v_t) and generate N paths from $t = 0$ to $t = T$. We then retain the last stock price from each stock price path and obtain the payoff of the European option at expiry, take the average over all stock price paths and discount back to time zero. So, in the case considered we would have respectively for a call option, $C(K)$ and a put option $P(K)$:

$$C(K) = e^{-rT} \frac{1}{N} \sum_{i=1}^{N} \max\left(0, S_T^{(i)} - K\right) \quad (4)$$

$$P(K) = e^{-rT} \frac{1}{N} \sum_{i=1}^{N} \max\left(0, K - S_T^{(i)}\right) \quad (5)$$

Where $S_T^{(i)}$ is the final price of the stock or index generated by the i -th path for $i = 1, \dots, N$. Such valuation requires the estimation of the characteristic parameters of the Heston dynamics: κ, θ, σ, v_0 and ρ (refer to paragraph 3).

There are two kinds of problems that arise when simulating the bivariate stochastic process (S_t, v_t). The first aspect is the slow convergence speed. The second, to be considered as a more serious problem, is given by the CIR (Cox-Ingersoll-Ross) type of process which describes the variance v_t over time. Considering how it was modeled, this dynamic leads a large number of numerical simulation schemes (including the most popular in the financial industry, namely Euler and Milstein) to generate negative values for v_t, even though the Feller condition, such that $2\kappa\theta > \sigma^2$, is respected. This occurs because such condition is valid for continuous-time CIR stochastic processes, while simulations, by nature, work in discrete time, approximating the dynamics defined on the continuum. The simplest and most direct way to manage the problem of having a negative variance even though Feller's condition is satisfied is to correct it instantaneously by systematically introducing an override every time this unwanted effect occurs. There are at least two ways to implement this:

- in the full truncation scheme, a negative value for v_t is set equal to zero. So v_t is replaced with $v_t^+ = \max(0, v_t)$ anywhere during the discretization process.

- in the reflection scheme, a negative value for v_t is reflected with $-v_t$. So v_t is replaced with $|v_t|$ anywhere during the discretization process.

The disadvantage of the full truncation scheme is that it creates a zero variance, which is an inaccurate representation of the real dynamics of an asset, which never shows a zero variance.

The disadvantage of the reflection schemes is that they make a large negative value assume a largely positive one. So, in other words, it would produce the bias of turning a low volatility into a high volatility.

The first problem related to the Heston model can be mitigated by taking into consideration variance reduction methodologies (Giribone and Ligato, 2013) or by implementing a Randomized Quasi Monte Carlo (Giribone and Ligato, 2014), since there are no distorting effects of convergence.

For the second aspect, certainly more crucial, it is necessary to resort to an adjustment of the integration scheme itself in the hope of improving the approximation of the continuous dynamics. So, another way to deal with negative simulated values of v_t is to design simulation schemes for variance that do not inherently produce negative values or so that the probability of running into such cases is very low. Most of the research focuses precisely on this aspect, namely that of simulating the variance process in the Heston model in the most accurate and stable way possible.

All simulation schemes for the Heston model contain the same basic steps. First, two independent random draws are made from a standard normal distribution. These variables are made dependent by applying the Cholesky decomposition. They are then multiplied by $\sqrt{dt}$ to make them a proxy for Brownian motion increments. The second step provides the updated value of the variance v_{t+dt} and the last step the updated value for the share or index, S_{t+dt}. This procedure, common to all the schemes that will be presented below, can therefore be summarized as follows:

Initialization: Assign the spot value to S_0 and the initial variance value to v_0.

Step 1. Generate two independent random variables Z_1 and Z_2 and define $Z_V = Z_1$ and $Z_S = \rho Z_V + \sqrt{1-\rho^2} Z_2$. Approximate the Brownian motion with $dW_{1,t} = Z_S\sqrt{dt}$ and $dW_{2,t} = Z_V\sqrt{dt}$.

Step 2. Get the updated value for v_{t+dt}

Step 3. Given v_{t+dt} calculate the updated value of S_{t+dt} and return to Step 1.

Let us note that, in accordance with the traditional use of the Cholesky decomposition for the simulation of correlated variables, the following statistical properties hold: $E[Z_V] = E[Z_S] = 0$ and $E[Z_V Z_S] = \rho E[Z_1^2] + \sqrt{1-\rho^2} E[Z_1 Z_2] = \rho$.

In the following subparagraphs, the most common integration schemes for (S_t, v_t) will be described assuming that the time grid is discretized using equally spaced time increments with a size equal to dt.

2.1) The Euler scheme

The easiest way to discretize the process represented in Eq. (3) is to adopt the traditional Euler scheme. This is equivalent to approximating integrals using the left-point rule (Rouah, 2013). The first integral is approximated as the product of the integrand at time t and the integration domain dt:

$$\int_t^{t+dt} \mu(X_u, u)du \approx \mu(X_t, t) \int_t^{t+dt} du = \mu(X_t, t)dt \quad (6)$$

The left-point rule is used since the value $\mu(X_t, t)$ is known at time t. The second integral is approximated as follows:

$$\int_t^{t+dt} \sigma(X_u, u)du \approx \sigma(X_t, t) \int_t^{t+dt} dW_u = \sigma(X_t, t)(W_{t+dt} - W_t) = \sigma(X_t, t)\sqrt{dt}Z, \quad (7)$$

since $W_{t+dt} - W_t$ and $\sqrt{dt}Z$ are identical in distribution, where Z is a standard normal variable. Thus, the discretization of Equation (3) according to the Euler method is:

$$X_{t+dt} = X_t + \mu(X_t, t)dt + \sigma(X_t, t)\sqrt{dt}Z \quad (8)$$

Considering the specific Heston model, the next step is to particularize Eq. (8) for the dynamics that regulate the variance and for the one that regulates the price.

The SDE for v_t in Eq. (1) rewritten in the form of Eq. (3) is:

$$v_{t+dt} = v_t + \int_t^{t+dt} \kappa(\theta - v_u)du + \int_t^{t+dt} \sigma\sqrt{v_u}dW_u \quad (9)$$

In accordance with Eq. (8), the Euler discretization approximates the integrals in (9) as:

$$\int_t^{t+dt} \kappa(\theta - v_u)du \approx \kappa(\theta - v_t)dt \quad (10)$$

$$\int_t^{t+dt} \sigma\sqrt{v_u}dW_{2,u} \approx \sigma\sqrt{v_t}(W_{t+dt} - W_t) = \sigma\sqrt{v_t}\sqrt{dt}Z_V \quad (11)$$

This implies that the Euler discretization for the variance is:

$$v_{t+dt} = v_t + \kappa(\theta - v_t)dt + \sigma\sqrt{v_t}\sqrt{dt}Z_V \quad (12)$$

From the theory related to stochastic processes of the CIR type, the probability of generating negative values for v_{t+dt} can be calculated, as follows:

$$\Pr(v_{t+dt} < 0) = \Phi\left(\frac{-(1-\kappa dt)v_t - \kappa\theta dt}{\sigma\sqrt{v_t}\sqrt{dt}}\right) \quad (13)$$

Where $\Phi(x)$ denotes the standard normal cumulative distribution function, evaluated at x. Therefore, since there is a non-zero probability of having a negative variance, the full truncation scheme or a reflection scheme should be applied to override any negative value generated during the simulation.

Regarding the simulation of the stock price or the index price, there are two common approaches: we can either directly simulate S_t or we can simulate $\ln S_t$ then calculate the exponential on the result obtained. The SDE for S_t in Eq. (3) can be expressed in integral form as:

$$S_{t+dt} = S_t + (r - q)\int_t^{t+dt} S_u du + \int_t^{t+dt} \sqrt{v_u}S_u dW_u \quad (14)$$

Applying Eq. (8), the Euler discretization approximates the integrals as follows:

$$\int_t^{t+dt} S_u du \approx S_t dt \quad (15)$$

$$\int_t^{t+dt} \sqrt{v_u}S_u dW_{1,u} \approx \sqrt{v_t}S_t(W_{t+dt} - W_t) = \sqrt{v_t}S_t\sqrt{dt}Z_S \quad (16)$$

Consequently, the discretization of the share price or index price is:

$$S_{t+dt} = S_t + (r - q)S_t dt + \sqrt{v_t}S_t\sqrt{dt}Z_S \quad (17)$$

To simulate the log stock price, we apply Itô's lemma to the first dynamic in Eq. (1). Then, $\ln S_t$ follows the SDE:

$$d\ln S_t = \left(r - q - \frac{1}{2}v_t\right)dt + \sqrt{v_t}dW_{1,t} \quad (18)$$

Or expressing it as integral:

$$\ln S_{t+dt} = \ln S_t + \int_t^{t+dt}\left(r - q - \frac{1}{2}v_u\right)du + \int_t^{t+dt}\sqrt{v_u}dW_{1,u} \quad (19)$$

The Euler discretization for the process $\ln S_t$ is:

$$\ln S_{t+dt} \approx \ln S_t + \left(r - q - \frac{1}{2}v_t\right)dt + \sqrt{v_t}\left(W_{1,t+dt} - W_{1,t}\right) = \ln S_t + \left(r - q - \frac{1}{2}v_t\right)dt + \sqrt{v_t}\sqrt{dt}Z_S \quad (20)$$

The Euler discretization for S_t is obtained applying the exponential to the various terms of the previous equation:

$$S_{t+dt} = S_t \exp\left[\left(r - q - \frac{1}{2}v_t\right)dt + \sqrt{v_t}\sqrt{dt}Z_S\right] \quad (21)$$

Again, in order to avoid the undesirable effect of obtaining negative variances, the introduction of the full truncation or reflection scheme is necessary, replacing v_t respectively with v_t^+ or $|v_t|$.

To implement the Euler simulation, initialization is performed setting S with the initial values of S_0 (or in the case of adoption of the logarithmic evolution process $x_0 = \ln S_0$) for the stock price and v_0 for the variance (in all cases). Given the values (S_t, v_t), v_{t+dt} is obtained from Eq. (12) and S_{t+dt} is obtained both from Eq. (17) and from Eq. (21) (for the case of logarithmic evolution of the asset price over time).

2.2) The Milstein scheme

The general case for the discretization of an SDE according to such a numerical integration scheme is described in Glasserman (2003) and Kloeden and Platen (1992). For the specific case observed, i.e., the Heston bivariate process, the coefficients of Eq. (2) do not depend directly on time t, but exclusively on X_t. For the sake of simplicity, we can thus assume that the stock price and the variance are driven by the SDE:

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t = \mu_t dt + \sigma_t dW_t \quad (22)$$

In integral form:

$$X_{t+dt} = X_t + \int_t^{t+dt} \mu_s ds + \int_t^{t+dt} \sigma_s dW_s \quad (23)$$

The idea behind the Milstein scheme is that the accuracy of the discretization can be increased by expanding the coefficients $\mu_t = \mu(X_t)$ and $\sigma_t = \sigma(X_t)$ through Itô's lemma. The coefficients will follow the following SDEs:

$$d\mu_t = \left(\mu_t' \mu_t + \frac{1}{2}\mu_t'' \sigma_t^2\right) dt + (\mu_t' \sigma_t) dW_t \quad (24)$$

$$d\sigma_t = \left(\sigma_t' \mu_t + \frac{1}{2}\sigma_t'' \sigma_t^2\right) dt + (\sigma_t' \sigma_t) dW_t \quad (25)$$

Where the single $'$and double quotes $''$ refer to the differentiation in X and where the derivatives in t are zero since μ_t and σ_t do not have a direct time dependence in the Heston model. The integral form of the coefficients at time s (with $t < s < t + dt$) is:

$$\mu_s = \mu_t + \int_t^s \left(\mu_u' \mu_u + \frac{1}{2}\mu_u'' \sigma_u^2\right) du + \int_t^s (\mu_u' \sigma_u) dW_u \quad (26)$$

$$\sigma_s = \sigma_t + \int_t^s \left(\sigma_u' \mu_u + \frac{1}{2}\sigma_u'' \sigma_u^2\right) du + \int_t^s (\sigma_u' \sigma_u) dW_u \quad (27)$$

Substituting μ_s and σ_s into the integrals in Eq. (23) we obtain:

$$X_{t+dt} = X_t + \int_t^{t+dt} \left[\mu_t + \int_t^s \left(\mu_u' \mu_u + \frac{1}{2}\mu_u'' \sigma_u^2\right) du + \int_t^s (\mu_u' \sigma_u) dW_u\right] ds +$$

$$+ \int_t^{t+dt} \left[\sigma_t + \int_t^s \left(\sigma_u' \mu_u + \frac{1}{2}\sigma_u'' \sigma_u^2\right) du + \int_t^s (\sigma_u' \sigma_u) dW_u\right] dW_s \quad (28)$$

The first order major differentials $dsdu = \mathcal{O}(dt^2)$ and $dsdW_u = \mathcal{O}(dt^{3/2})$ are here neglected. The term $dW_u dW_s$ is retained since it is of the first order, $\mathcal{O}(dt)$. Considering such assumptions, Eq. (28) simplifies to:

$$X_{t+dt} = X_t + \mu_t \int_t^{t+dt} ds + +\sigma_t \int_t^{t+dt} dW_s + \int_t^{t+dt} \int_t^s (\sigma_u' \sigma_u)\, dW_u dW_s \quad (29)$$

We apply the Euler discretization to the last term of Eq. (29) and we obtain:

$$\int_t^{t+dt} \int_t^s (\sigma_u' \sigma_u)\, dW_u dW_s \approx \sigma_t' \sigma_t \int_t^{t+dt} \int_t^s dW_u dW_s = \sigma_t' \sigma_t \int_t^{t+dt} (W_s - W_t) dW_s = \sigma_t' \sigma_t \left[\int_t^{t+dt} W_s dW_s - W_t W_{t+dt} + W_t^2\right] \quad (30)$$

To solve the remaining integral in Eq. (30), we define $dY_t = W_t dW_t$. Using Itô's Lemma, it is easy to prove that Y_t has the following solution: $Y_t = \frac{1}{2}W_t^2 - \frac{1}{2}t$. Therefore: $\frac{\partial Y}{\partial t} = -\frac{1}{2}, \frac{\partial Y}{\partial W} = W$ and $\frac{\partial^2 Y}{\partial W^2} = 1$:

$$dY_t = \left(-\frac{1}{2} + 0 + \frac{1}{2} \times 1 \times 1\right) dt + (W_t \times 1) dW_t = W_t dW_t \quad (31)$$

Using this result, we can write:

$$\int_t^{t+dt} W_s dW_s = Y_{t+dt} - Y_t = \frac{1}{2}W_{t+dt}^2 - \frac{1}{2}W_t^2 - \frac{1}{2}dt \quad (32)$$

Substituting this last equation into Eq. (30) we obtain:

$$\int_t^{t+dt} \int_t^s (\sigma_u' \sigma_u)\, dW_u dW_s \approx \frac{1}{2}\sigma_t' \sigma_t [(W_{t+dt} - W_t)^2 - dt] = \frac{1}{2}\sigma_t' \sigma_t [(\Delta W_t)^2 - dt] \quad (33)$$

Where $\Delta W_t = W_{t+dt} - W_t$, which is equal in distribution to $\sqrt{dt}Z$ with Z distributed as a standard normal. Combining Eq. (29) and Eq. (33), the general form for the Milstein discretization is therefore:

$$X_{t+dt} = X_t + \mu_t dt + \sigma_t \sqrt{dt} Z + \frac{1}{2}\sigma_t' \sigma_t dt (Z^2 - 1) \quad (34)$$

Thus, the Milstein discretization for dX_t expressed in Eq. (34) is identical to the Euler one in Eq. (8), with the exception of the added term $\frac{1}{2}\sigma_t' \sigma_t dt (Z^2 - 1)$ which allows to improve the accuracy of the discretization scheme compared to the standard one.

Reconsidering the Heston model in Eq. (1), this last integration scheme can be applied both for the processes of S_t (or $\ln S_t$) and the processes of v_t.

The coefficients for the variance process are $\mu(v_t) = \kappa(\theta - v_t)$ and $\sigma(v_t) = \sigma\sqrt{v_t}$, substituting them in the general expression (34) we obtain:

$$v_{t+dt} = v_t + \kappa(\theta - v_t)dt + \sigma\sqrt{v_t}\sqrt{dt} Z_V + \frac{1}{4}\sigma^2 dt (Z_V^2 - 1) \quad (35)$$

Which can be rewritten as:

$$v_{t+dt} = \left(\sqrt{v_t} + \frac{1}{2}\sigma\sqrt{dt} Z_V\right)^2 + \kappa(\theta - v_t)dt - \frac{1}{4}\sigma^2 dt \quad (36)$$

Although the Milstein discretization scheme for the variance stochastic process produces far fewer negative values compared to the basic Euler scheme (Rouah, 2013), it is still necessary to implement the full truncation scheme or the reflection scheme to Eq. (35) and Eq. (36).

The coefficients for the stock price or the index price process are $\mu(v_t) = (r - q)S_t$ and $\sigma(S_t) = \sqrt{v_t} S_t$, substituting them into the general expression, Eq. (34) becomes:

$$S_{t+dt} = S_t + (r-q)S_t dt + \sqrt{v_t}\sqrt{dt} S_t Z_S + \frac{1}{2} v_t S_t dt (Z_S^2 - 1) \quad (37)$$

We can also discretize the log-stock process, which, according to Itô's lemma, will follow the following dynamics:

$$d \ln S_t = \left(r - q - \frac{1}{2} v_t\right) dt + \sqrt{v_t} dW_{1,t} \quad (38)$$

The coefficients are $\mu(S_t) = r - q - \frac{1}{2} v_t$ and $\sigma(S_t) = \sqrt{v_t}$ so that $\mu_t' = \sigma_t' = 0$. Since v_t is known at time t, we can treat it as a constant within the definition of the coefficients. Applying Eq. (34), we reach the following result:

$$\ln S_{t+dt} = \ln S_t + \left(r - q - \frac{1}{2} v_t\right) dt + \sqrt{v_t}\sqrt{dt} Z_s \quad (39)$$

which is identical to Eq. (20) obtained applying the simple Euler discretization. Thus, the Milstein discretization of $\ln S_t$ in the Heston model does not produce a more accurate approximation than the basic method.

Similarly to what has already been discussed, the price of the share or the index is directly obtained applying the exponential function to all the terms of Eq. (39). The adoption of the full truncation or the reflection scheme is also necessary for Eq. (37) or Eq. (39).

To implement the Milstein simulation, we initialize the value of the stock price and the value of the variance, respectively with S_0 and v_0. Given the simulated values (S_t, v_t), v_{t+dt} is obtained from Eq. (35) and S_{t+dt} is obtained from Eq. (37) or, alternatively, from Eq. (39).

In Eq. (34) for the discretization of dX_t, the coefficients $\mu_t = \mu(X_t)$ and $\sigma_t = \sigma(X_t)$ are respectively the drift and the volatility of the process for X_t and they are functions of X_t itself. In the implicit version of the Milstein scheme, the coefficient of the drift μ_t is expressed as a function of X_{t+dt}. Consequently, such value is known only implicitly and not explicitly, as in the previous case which depended on X_t and Eq. (34) becomes:

$$X_{t+dt} = X_t + \mu_{t+dt} dt + \sigma_t \sqrt{dt} Z + \frac{1}{2} \sigma_t' \sigma_t dt (Z^2 - 1) \quad (40)$$

Where $\mu_{t+dt} = \mu(X_{t+dt})$. It is also possible to interpolate between the Milstein implicit-explicit schemes, calculating a weighted average between μ_t and μ_{t+dt}. The weighted implicit-explicit Milstein scheme is therefore:

$$X_{t+dt} = X_t + [\alpha \mu_t + (1-\alpha)\mu_{t+dt}] dt + \sigma_t \sqrt{dt} Z + \frac{1}{2} \sigma_t' \sigma_t dt (Z^2 - 1) \quad (41)$$

Where $\alpha \in (0,1)$ represents the assigned weight. The explicit Milstein can be obtained as a degenerate case by setting $\alpha = 1$, just as the implicit Milstein can be obtained by setting $\alpha = 0$.

To apply the implicit Milstein scheme to the Heston model, in Eq. (35) we replace the term $\kappa(\theta - v_t)dt$ with $\kappa(\theta - v_{t+dt})dt$. We then bring $\kappa v_{t+dt} dt$ over to the left-hand side of the resulting equation, and we divide by $1 + \kappa dt$ to obtain:

$$v_{t+dt} = \frac{v_t + \kappa\theta dt + \sigma\sqrt{v_t}\sqrt{dt} Z_V + \frac{1}{4}\sigma^2 dt (Z_V^2 - 1)}{1 + \kappa dt} \quad (42)$$

The same steps can also be performed for the case of the weighted scheme (Eq. 41) to obtain:

$$v_{t+dt} = \frac{v_t + \kappa(\theta - \alpha v_t) dt + \sigma\sqrt{v_t}\sqrt{dt} Z_V + \frac{1}{4}\sigma^2 dt (Z_V^2 - 1)}{1 + (1-\alpha)\kappa dt} \quad (43)$$

2.3) The Transformed Volatility scheme

One way to avoid negative variances is to simulate volatility rather than variance, and then square the result. From Itô's lemma, the volatility $\omega_t = \sqrt{v_t}$ follows the following stochastic process (Rouah, 2013):

$$d\omega_t = \frac{\kappa}{2}\left[\left(\theta - \frac{\sigma^2}{4\kappa}\right)\frac{1}{\omega_t} - \omega_t\right] dt + \frac{1}{2}\sigma dW_{1,t} \quad (44)$$

The Euler discretization for Eq. (44) is:

$$\omega_{t+dt} = \omega_t + \frac{\kappa}{2}\left[\left(\theta - \frac{\sigma^2}{4\kappa}\right)\frac{1}{\omega_t} - \omega_t\right] dt + \frac{1}{2}\sigma\sqrt{dt} Z_V \quad (45)$$

While the Euler discretization of the log stock price produces:

$$S_{t+dt} = S_t \exp\left[\left(r - q - \frac{1}{2}\omega_t^2\right) dt + \omega_t \sqrt{dt} Z_S\right] \quad (46)$$

Zhu (2010) has shown that the Euler discretization of the volatility ω_t avoids the negative variances, but it has the disadvantage that the mean level $\theta_\omega = \frac{\theta - \frac{\sigma^2}{4\kappa}}{\omega_t}$ in Eq. (44) is stochastic due to the term $\frac{1}{\omega_t}$.

This could cause low simulation performance. The transformed volatility scheme proposed by Zhu (2010) applies a robust approximation for θ_ω which allows to rectify this problem. His transformed process for volatility is:

$$d\omega_t = \frac{\kappa}{2}[\theta^* - \omega_t] dt + \frac{\sigma}{2} dW_{2,t} \quad (47)$$

Which is characterized by having a mean reversion speed equal to $\kappa/2$ and a variance volatility equal to $\sigma/2$.

The mean reversion level θ_t^* is equal to:

$$\theta_t^* = \frac{\beta - \omega_t \exp(-\kappa dt/2)}{1 - \exp(-\kappa dt/2)} \quad (48)$$

Where:

$$\beta = \sqrt{[E(v_{t+dt}) - \mathrm{Var}(\omega_{t+dt})]^+} = \sqrt{\left[\theta + (v_t - \theta)\exp(-\kappa dt) - \frac{\sigma^2}{4\kappa}(1 - \exp(-\kappa dt))\right]^+} \quad (49)$$

Note that parameter β is equal to zero when $E(v_{t+dt}) < \mathrm{Var}(\omega_{t+dt})$, while the mean reversion level θ_t^* depends on the value of ω_t.

The Euler discretization for $d\omega_t$ in Eq. (47) generates:

$$\omega_{t+dt} = \omega_t + \frac{\kappa}{2}[\Theta^* - \omega_t]dt + \frac{\sigma}{2}\sqrt{dt}Z_V \quad (50)$$

The same procedure can be used for conducting the simulations as discussed in the previously presented discretization schemes, with the sole exception that the starting point for the initial volatility should be set to $\sqrt{v_0}$.

2.4) The Balanced Implicit scheme

This scheme is able to preserve positivity in the stochastic process associated with the variance. It is defined in Platen and Heath (2009) and in Kahl and Jäckel (2006) as:

$$v_{t+dt} = v_t + \mu_t dt + \sigma_t \Delta W_t + (v_t - v_{t+dt})C(v_t) \quad (51)$$

Where:

$$C(v_t) = c^0(v_t)dt + c^1(v_t)|\Delta W_t| \quad (52)$$

With $c^0(v_t) = \kappa$ and $c^1(v_t) = \sigma/\sqrt{v_t}$. The Balanced Implicit scheme for the Heston model is therefore:

$$v_{t+dt} = v_t + \kappa(\theta - v_t)dt + \sigma\sqrt{v_t}\sqrt{dt}Z_V + (v_t - v_{t+dt})C(v_t) = \frac{v_t[1 + C(v_t)] + \kappa(\theta - v_t)dt + \sigma\sqrt{v_t}\sqrt{dt}Z_V}{1 + C(v_t)} \quad (53)$$

With:

$$C(v_t) = \kappa dt + \frac{\sigma\sqrt{dt}|Z_V|}{\sqrt{v_t}} \quad (54)$$

Since the variance is always guaranteed to be positive, reflection and full truncation schemes are unnecessary. Unfortunately, as shown by Kahl and Jäckel (2006), the convergence of this scheme is not always optimal.

2.5) The Pathwise Adapted Linearization Quadratic

Another scheme for variance discretization is the Path Adapted Linearization Quadratic introduced by Kahl and Jäckel (2006). These authors demonstrated a faster convergence compared to the previous one, especially for small values of σ. The discretization scheme is given by:

$$v_{t+dt} = v_t + \left[\kappa(\tilde{\theta} - v_t) + \sigma\beta_n\sqrt{v_t}\right] + \left(1 + \frac{\sigma\beta_n - 2\kappa\sqrt{v_t}}{4\sqrt{v_t}}dt\right)dt \quad (55)$$

Where $\tilde{\theta} = \theta - \sigma^2/(4\kappa)$ and where $\beta_n = \frac{Z_V}{\sqrt{dt}}$. For high values of σ it could introduce potential instabilities (Rouah, 2013).

2.6) The Kahl-Jäckel IJK Scheme

This scheme was also proposed by Kahl and Jäckel (2006) and it consists in simulating v_t with the implicit Milstein scheme, according to Eq. (42) and simulating the $\ln S_t$ with the IJK discretization:

$$\ln S_{t+dt} = \ln S_t + \left(r - q - \frac{v_t + v_{t+dt}}{4}\right)dt + \rho\sqrt{v_t dt}Z_V + \frac{1}{2}\left(\sqrt{v_t} + \sqrt{v_{t+dt}}\right)(Z_S - \rho Z_V)\sqrt{dt} + \frac{\rho\sigma dt}{2}(Z_V^2 - 1) \quad (56)$$

Since this scheme can produce negative values for variance, it is necessary to implement it in conjunction with the full truncation or the reflection scheme. For details on the derivation of the scheme, please refer to the above cited paper.

2.7) The Moment Matching scheme

Andersen and Brotherton-Ratcliffe (2005) proposed a moment-matched discretization scheme which generates only positive variances. This technique produces a variance which is distributed according to a log-normal, so a reasonable choice of parameterization is to match the first two moments of a log-normal discretization process.

The numerical integration scheme can therefore be expressed in the following form:

$$v_{t+dt} = [\theta + (v_t - \theta)\exp(-\kappa dt)]\exp\left(-\frac{1}{2}\Gamma_t^2 + \Gamma_t Z_V\right) \quad (57)$$

Where:

$$\Gamma_t = \ln\left(1 + \frac{\sigma^2 v_t[1 - \exp(-2\kappa dt)]}{2\kappa[\theta + (v_t - \theta)\exp(-\kappa dt)]^2}\right) \quad (58)$$

Following the part of the study dedicated to the theoretical description of the most popular numerical schemes associated with the Heston bivariate stochastic model, the ensuing sections of the paper focus on their algorithmic implementation and, consequently, on the scientific evidence found. In particular, the next paragraph examines the process of estimating the parameters of the dynamics which typically takes place using call and put options listed on the market and written on the same certificate underlying. This calibration occurs by comparing the market prices with the theoretical prices computed using the closed formula for vanilla options derived from Heston (1993) and successively improved by (Albrecher et al., 2007).

Once the five characteristic parameters have been estimated, we move on to the programming of the discussed numerical schemes: Euler [E], Explicit Milstein [M], Implicit Milstein [IM], Weighted Implicit-Explicit Milstein [WM], Transformed Volatility [TV], Balanced Implicit [B], Pathwise Adapted Linearization Quadratic [PW], Kahl-Jackel [IJK] and Moment Matching [MM]. Also in section 3, the code and the absence of bias are verified through the comparison between the output of the Monte Carlo simulator and the analytical formula provided by Heston (1993) for the case of standard optionality (European calls/puts). Once we are confident on the validity of the implemented model, we proceed to the last part of the paper and to the pricing of two investment certificates having a highly non-linear optionality and in line with the characteristics that are most requested by customers (Acepi associates' primary market 2023 Q1). In this case there are no analytical pricing formulas for certain types of features such as autocallability or the memory effect associated with the payment of additional amounts. The comparison therefore takes place exclusively between the various numerical schemes and we estimate how many times it was necessary to adjust the variance for each of the integration schemes.

3) Implementation of the numerical integration schemes for the Heston model

The Heston model implies that the price drift and volatility of a security follow certain laws according to 5 parameters: $V_0, \theta, k, \sigma, \rho$. These parameters cannot be directly observed on the market therefore they must be estimated in order to enter them into the Monte Carlo pricing engine. For the calibration, we start from the implied volatility surface of the FTSE MIB and take the volatilities of the options traded on the market, then all the parameters are estimated together using a least squares minimization (Mrázek & Pospíšil, 2017). It is also possible to assign weights, for example, giving importance to the volatilities deriving from the most traded options on the market (the most liquid ones). Such implied surface has strikes ranging from 80% to 120% in terms of moneyness while the maturities range from 1 month to 7 years. The calibration is presented as a five-dimension minimization problem where we try to minimize the least squares of the differences between the volatilities obtained from the model and those observed on the market. Therefore, defining the implied volatility of an option as $I(V_i)$, the problem is as follows:

$$\min \sum_{i=1}^{n} (I(V_i^{model}(S, t_i, K_i, \bar{\phi})) - I(V_i^{market}(S, t_i, K_i))^2 \quad (59)$$

With $\overline{\phi} = (V_0, \theta, \kappa, \sigma, \rho)$ under the following conditions: $V_t \geq 0,\ \theta \geq 0, k \geq 0, \sigma \geq 0$ and $-1 \leq \rho \leq +1$

The closed formula for pricing a European Call Option that pays a continuous dividend within the Heston model pricing framework is (Heston, 1993):

$$C(S_t V_t, t, T) = S_t P_1 - K e^{-r(T-t)} P_2 \quad (60)$$

Where:

$$P_j(x, V_t, T, K) = \frac{1}{2} + \frac{1}{\pi} \int_0^{\infty} \mathrm{Re}\left(\frac{e^{-i\phi \ln(K)} f_j(x, V_t, T, \phi)}{i\phi}\right) d\phi \quad (61)$$

$$x = \ln(S_t) \quad (62)$$

$$f_j(x, V_t, T, \phi) = \exp\{C(T-t, \phi) + D(T-t, \phi)V_t + i\phi x\} \quad (63)$$

$$C(T-t, \phi) = r\phi i(r-q)\frac{a}{\sigma^2}\left[\left(b_j - \rho\sigma\phi i + d\right)\tau - 2\ln\left(\frac{1-ge^{d\tau}}{1-g}\right)\right] \quad (64)$$

$$D(T-t, \phi) = \frac{b_j - \rho\sigma\phi i + d}{\sigma^2}\left(\frac{1-e^{d\tau}}{1-ge^{d\tau}}\right) \quad (65)$$

$$g = \frac{b_j - \rho\sigma\phi i + d}{b_j - \rho\sigma\phi i - d} \quad (66)$$

$$d = \sqrt{(\rho\sigma\phi i - b_j)^2 - \sigma^2(2u_j\phi i - \phi^2)} \quad (67)$$

For $j = 1, 2$ where: $u_1 = \frac{1}{2},\ u_2 = -\frac{1}{2},\ a = k\theta,\ b_1 = k - \rho\sigma,\ b_2 = k$

In the formulas, i represents the imaginary unit.

In the cases considered in this study, let us consider the market data of the implied Black-Scholes volatilities associated with the FTSE MIB index as of 16 May 2023, whose surface $\sigma_{IMPL}(K, T)$ is shown in Figure 1.

In order to reach a reliable estimation of the parameters associated with the Heston bivariate stochastic model, we should only consider the most liquid European options provided by the Equity/Index markets.

We have consequently applied a filtering of the options in order to select only the most relevant for the calibration process.

The first criterion applied is to consider, in case of multiple contributions, the prices characterized by higher traded volumes.

The second criterion deals with the bid-ask spread: we exclude from the calibration process all the options for which the relative distance between the bid and the ask contribution is greater than a threshold α: $\frac{P^{ASK} - P^{BID}}{P^{BID}} > \alpha$. Typical values for α range between 50% and 70%.

The third criterion regards the so-called "penny option filter". We exclude all the options characterized by having a Bid price which is less than a percentage threshold (β) of the distance of the quoted strike (ΔK): $P^{BID} < \beta \cdot \Delta K$. Typical values for β range between 20% and 30%.

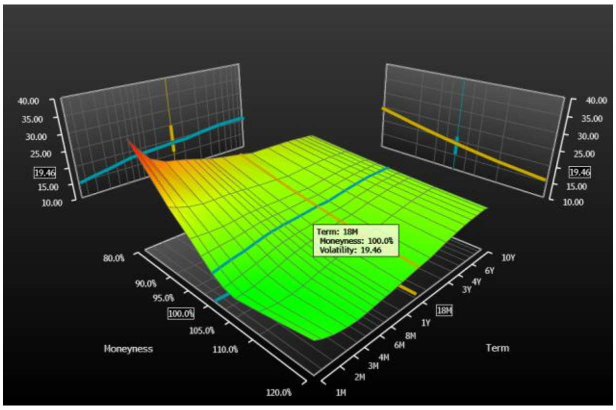

Expiry	Exp Date	Impl Fwd	Risk Free	Impl Dvd	Impl (Yld)
1M	16 Jun 2023	26984.17	3.658%	298.344	12.924%
2M	16 Jul 2023	27026.25	3.658%	337.225	7.424%
3M	16 Aug 2023	27009.70	3.658%	436.887	6.377%
6M	16 Nov 2023	27138.48	3.658%	555.909	4.057%
9M	16 Feb 2024	27168.37	3.811%	801.833	3.901%
1Y	16 May 2024	26976.30	3.785%	1226.620	4.501%
18M	16 Nov 2024	26943.06	3.628%	1688.351	4.122%
2Y	16 May 2025	26630.05	3.434%	2337.476	4.294%
3Y	16 May 2026	26225.53	3.186%	3364.202	4.122%
4Y	16 May 2027	25829.47	3.064%	4348.526	3.997%
5Y	16 May 2028	25456.61	3.008%	5298.472	3.894%
7Y	16 May 2030	24774.83	2.966%	7070.005	3.713%
10Y	16 May 2033	23973.45	2.992%	9425.802	3.465%

Figure 1: Implied Volatility Surface, Continuous Dividend Yield and Risk Free used for testing the numerical integration schemes– Source: Bloomberg®

Filtering the option market data based on the above criteria, 63 options are used to estimate the five parameters characterizing the Heston dynamics. The values thus obtained applying the minimization reported in Eq. 59 with the market data shown in Figure 1 are as follows:

Initial Variance: $v_0 = 2.923\%$

Variance reversion speed: $\kappa = 2.0517$

Variance reversion level: $\Theta = 0.05572$

Volatility of variance: $\sigma = 0.356267$

Correlation between the two Brownian motions: $\rho = -0.738102$

Using these parameters, the exact analytical formula for the pricing of a call option was implemented in Matlab with the aim of valuing call options with strikes in the moneyness range of 80% - 120% with step equal to 5% (spot equal to the closing value of the index on the calibration date, $S = 27198.9$) and a time to maturity in the range of 3 months – 5 years, with quarterly intervals. The pricing surface thus calculated is shown in Figure 2.

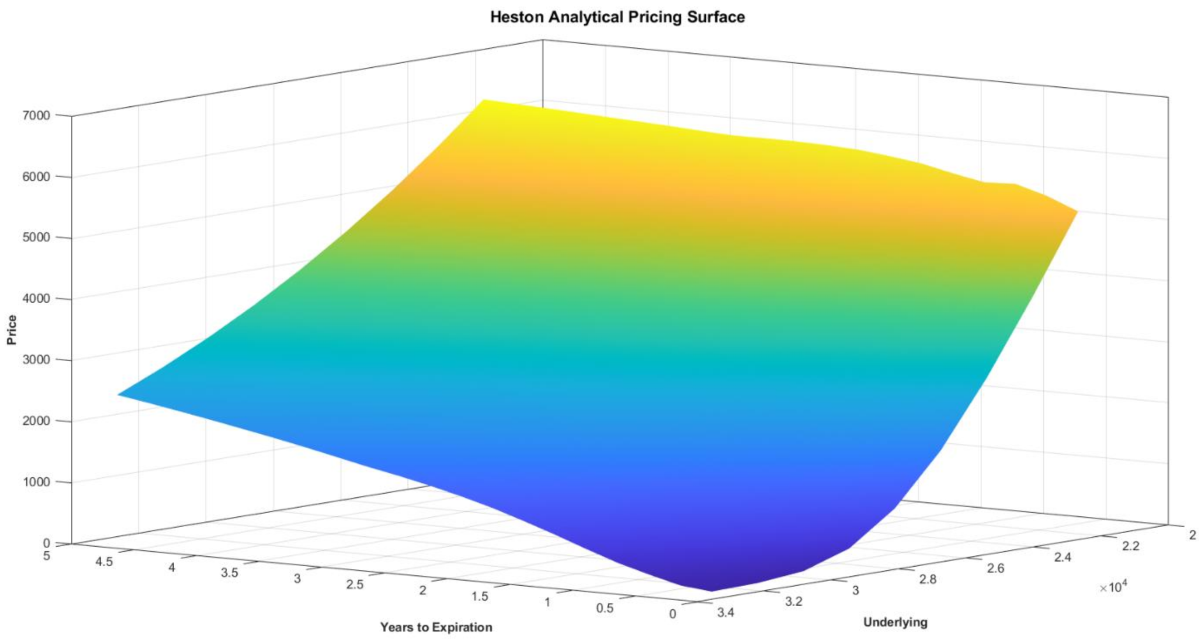

Figure 2: Pricing Surface of the call options written on the FTSE MIB index, priced with the Heston analytical formula

The next step was the programming of all the numerical schemes presented in paragraph 2. The comparison between the results obtained from the Monte Carlo method and the analytical formulas allows us to understand on the one hand the robustness of the discretization schemes (absence of bias), and on the other hand, the potential override of negative variances which unfortunately characterize the Euler traditional approach.

In fact, setting a number of paths equal to 10,000 and a constant discretization interval dt of one day, a non-negligible number of variances are corrected with the full truncation scheme (see Figure 3).

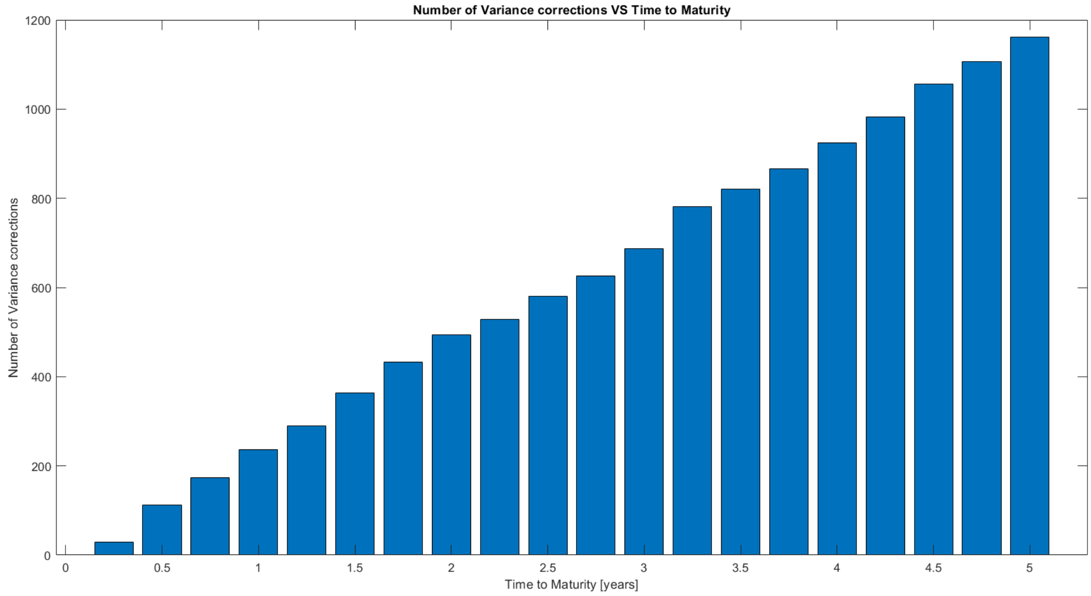

Figure 3: Number of variance corrections in the traditional Euler scheme

The discretization error introduced by the different schemes in the Monte Carlo remained stable and it is consistent with the expectations with respect to the chosen time step.

Furthermore, it should be noted that in all schemes other than Euler the event of obtaining a negative variance never occurred.

4) Numerical simulations for different typologies of certificates

Using market data from 16 May 2023 on the FTSE MIB Index (spot, dividend yield and implied volatility, shown in Figure 1), the interest rates term structure (Figure 1) and the five parameters of the Heston dynamics estimated in paragraph 3, we proceed with the valuation, according to the different numerical schemes described in paragraph 2, of two investment certificates having the characteristics most requested by investors, according to the ACEPI statistics for the first quarter of 2023.

Both certificates are characterized by the same underlying (FTSE MIB), the same maturity date of 21 September 2026 and they both pay coupons conditional on exceeding a barrier level equal to 70% of the initial reference value of the index, which we assumed to be the closing value of the index as of 20 September 2021 (25,048.26).

The value of the conditional amount was set at 3.55% and characterized by an annual frequency, so the four future Coupon Valuation Dates are: 20 September 2023, 20 September 2024, 22 September 2025 and 21 September 2026. We assumed for simplicity of account that the payment dates of the coupon coincide with the valuation dates.

Both certificates analyzed have conditionally protected capital at maturity: if the underlying index at maturity (September 21, 2026) is above the barrier level, set as the coupon trigger equal to 17,533.78, then the entire amount invested in the certificate will be returned; otherwise the negative performance of the underlying will be paid (thus, the redemption is characterized by a traditional pay-off of a short position in a barrier put option with a "European" observation of the barrier level (Giribone and Revetria, 2021), i.e. only at maturity).

The two certificates, on the other hand, differ for two aspects: the first structured product is characterized by standard digital coupons and no autocallability, while the second envisages both the memory effect and the possibility of early automatic callability if, on the Coupon Valuation Dates, the underlying index exceeds the initial reference value (i.e., 25,048.26).

It should be remembered that by memory effect we mean an interest payment that is carried over to the next observation dates if the product, at a given valuation date, fails to meet the coupon payment requirements as defined in the structure.

However, if payment requirements are met at a certain valuation date, all coupons that have not previously been paid will fall due for payment at such date.

We can express the rule for the payment of the t-th conditional amount in accordance with the following formula:

$$Nominal\ Value \times [t - th\ Additional\ Amount[\%] + Memory\ Additional\ Amount\ [\%] \times (t - k - 1)] \quad (68)$$

Where k can assume values between 0 and the number of coupon payment dates and indicates the value of t corresponding to the last conditional additional amount event.

The pay-offs of the two certificates analyzed in this study are quite popular and their features are classified in accordance with the EUSIPA (European Structured Investment Products Association) derivative map under the identification code 1260 - Express Certificates with additional coupon amount. ACEPI has adopted the same European classification typology.

Given that these two structured products are characterized by standard financial features, we suggest consulting the websites of these two associations for further information on the payoff mechanics. In order to value the two certificates, 500 replications of 10,000 simulations each have been implemented on the underlying index, with a daily time discretization step, until the maturity date.

The adjustments for the negative variance, applying the full truncation scheme, were only necessary in the case of the Euler integration method. The distribution of the total number of times for which the full truncation scheme was invoked throughout the 10,000 simulated paths for the 500 replications is shown in Figure 4. As regards the determination of the fair value and its measure of dispersion for the two certificates, all the calculated outputs resulted in line with our expectations given that we have replicated these calculations using other pricing modules, such as the library provided by Bloomberg ® - DLIB. All the results shown in Table 1 have been estimated using the same seed values for the generation of the random numbers stream for each replication. We consider this part very important in order to fairly compare the Monte Carlo outputs across the different numerical schemes. The results for the first and the second certificate are obtained using 10,000 paths and a discretization step equal to one day. Only a slight discrepancy is shown in correspondence with the approximation introduced by Moment Matching.

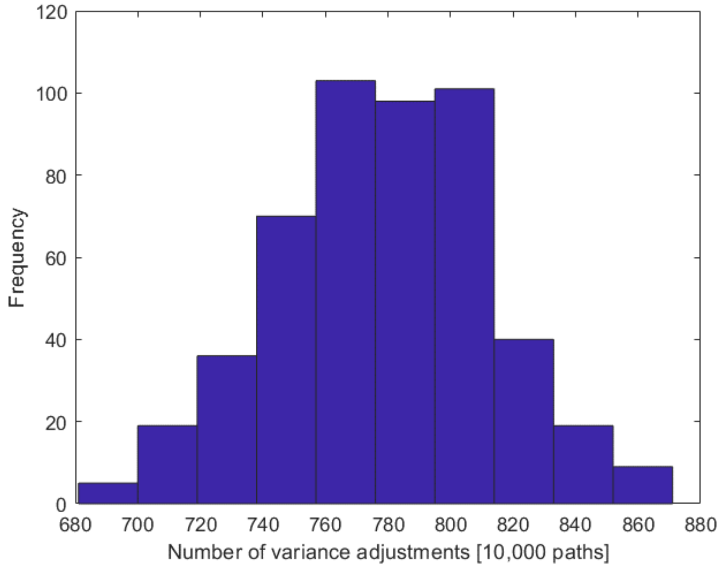

Figure 4: Distribution of variance adjustments in the traditional Euler numerical integration scheme in the example of investment certificates pricing

Numerical Integration Scheme	Fair Value Cert 1	St. Dev. Cert. 1	Fair Value Cert. 2	St. Dev. Cert. 2
Euler [E]	925.2473	2.1653	970.7499	1.4981
Explicit Milstein [M]	925.2754	2.1724	970.7566	1.4985
Implicit Milstein [IM]	925.3123	2.1683	970.7601	1.5012
Weighted Implicit-Explicit Milstein [WM]	925.2952	2.1706	970.7576	1.5008
Transformed Volatility [TV]	925.2381	2.1741	970.7435	1.4993
Balanced Implicit [B]	925.2437	2.1528	970.9811	1.4925
Pathwise Adapted Linearization Quadratic [PW]	925.2174	2.1716	970.7295	1.5011
Kahl-Jackel [IJK]	925.2173	2.1575	970.7443	1.4961
Moment Matching [MM]	926.3386	2.0658	967.6476	1.5211

Table 1: Pricing of the two certificates with different numerical integration schemes. All the measures are expressed in Euro and assuming a nominal value equal to 1,000

Further experiments have been conducted relaxing the constraints imposed on the seeds of the random number generator, with the aim of having the certainty that the results have not been compromised by the specific choice of the random source.

We have made ten replications for a number of simulations that varies from 1,000 to 20,000 with a step equal to 1,000 for both certificates. The alternative numerical schemes that we have implemented display similar performances (see Figures 5 and 6). This similarity can be explained considering that the adjustments for negative variances have never been used for all the simulations, as a result the positive effects on the fair values estimation introduced by the implementation of these schemes is clear.

In order to understand if the performances of the alternative integration methods also work well in a stressed scenario, we have replicated the same pricing experiments using the shocks of the European Banking Authority (EBA) 2023 available on the institution's website (https://www.eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing).

In particular we have applied the following shocks:

- a relative change in the spot level equal to -58.2%.

- an absolute change in the yield of interest rates term structure equal to: 152 bps up to 1 year and 167 bps from 1 year to 5 years. Zero-rates in the periods have been linearly interpolated.

- an absolute change in the credit spread of the issuer equal to + 171 bps, as a result we move from a risk-free to a risk-adjusted evaluation.

- a relative change in the Equity volatility of 268 basis points applied to the Heston v_0 parameter.

As shown in Figures 7 and 8, the price collapses and the severe shocks do not compromise the performances of the alternative methods in terms of the number of negative variance adjustments and, consequently, the standard deviations of the ten replications expressed in function of the number of paths maintain the similarity.

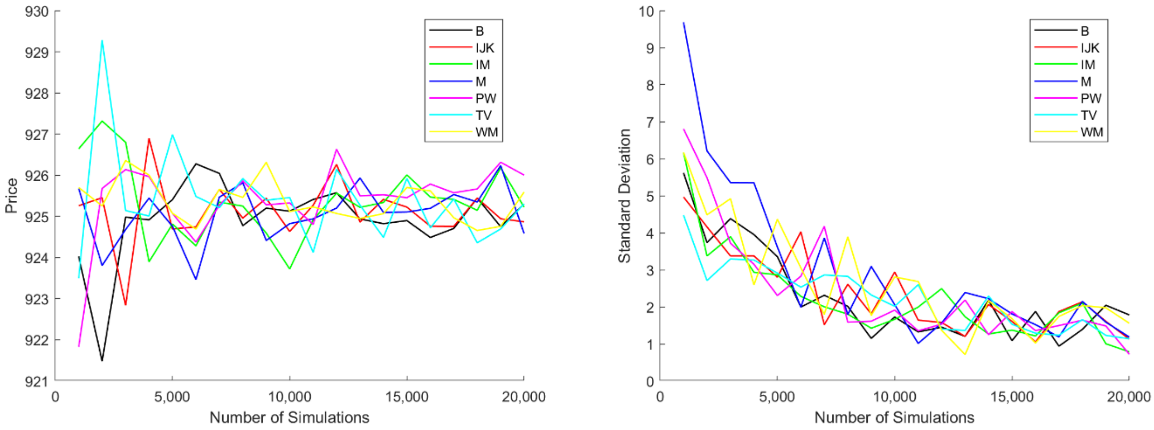

Figure 5: Prices and Standard Deviations of the Monte Carlo Outputs for the first Certificate as the number of paths varies – Base Scenario

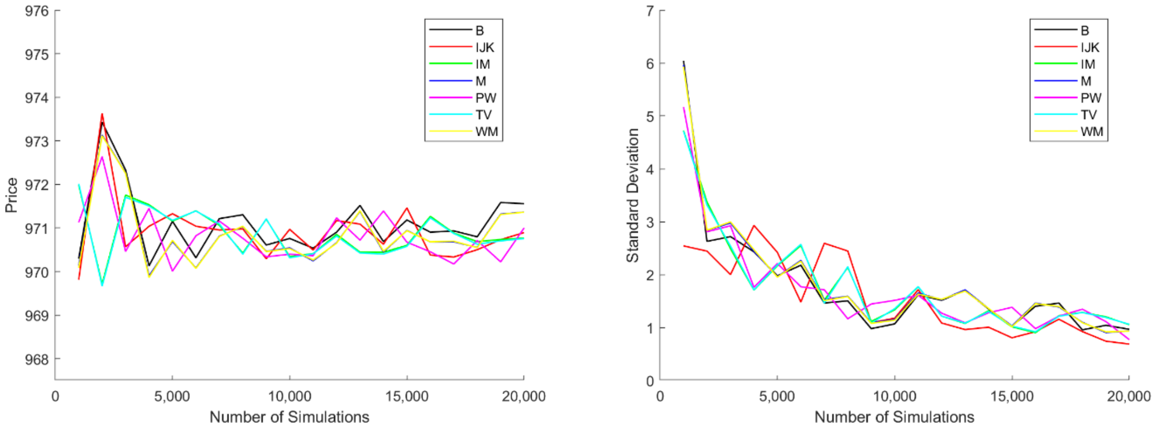

Figure 6: Prices and Standard Deviations of the Outputs for the second Certificate as the number of paths varies – Base Scenario

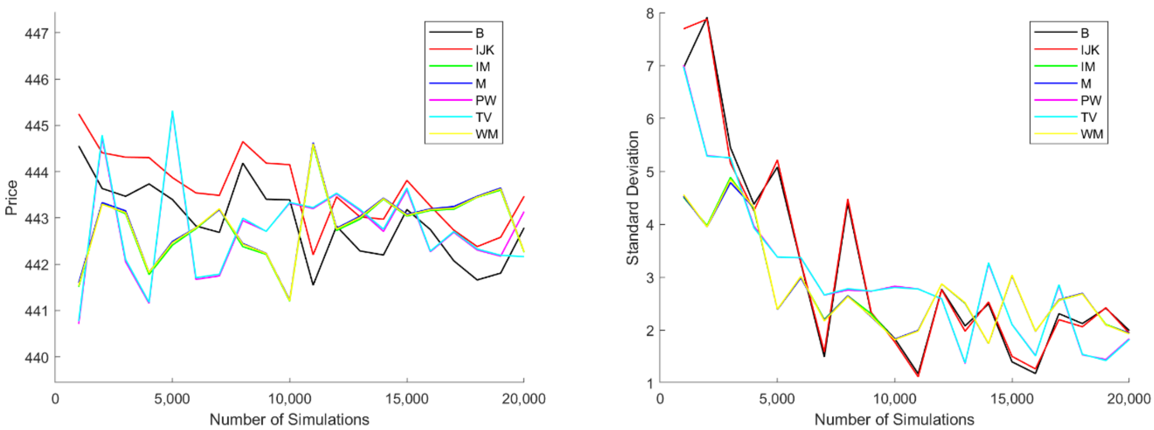

Figure 7: Prices and Standard Deviations of the Monte Carlo Outputs for the first Certificate as the number of paths varies – Stressed Scenario

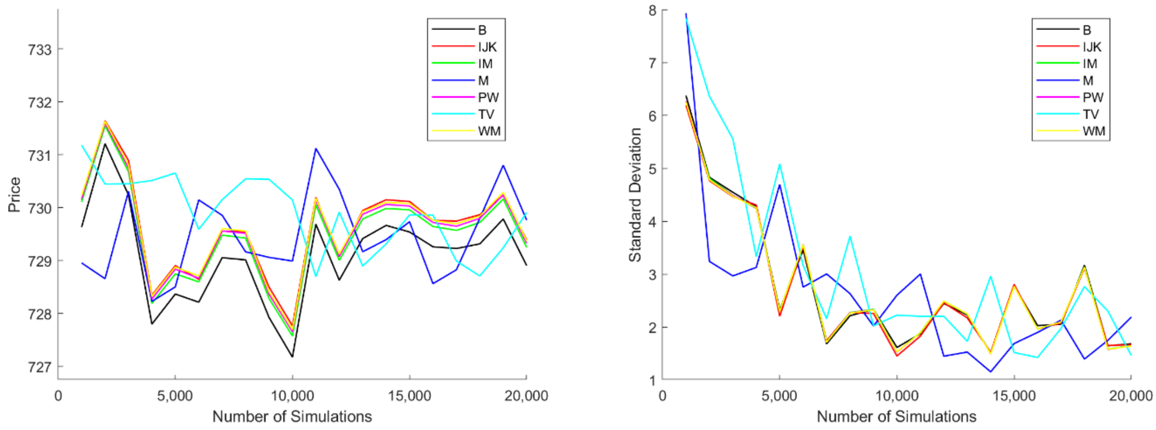

Figure 8: Prices and Standard Deviations of the Monte Carlo Outputs for the second Certificate as the number of paths varies – Stressed Scenario

The Heston model is one of the most widespread pricing frameworks in the financial industry as it allows a good trade-off between the complexity of the model, capable of efficiently representing the dynamics of the underlying price together with the volatility, and the possibility of estimating its characteristic parameters starting from the Black-Scholes log-normal volatilities through an analytical formula. It should be remembered that if we remove the stochastic contribution of the variance, we obtain the traditional dynamics of the Brownian geometric motion ruling the Black-Scholes-Merton pricing framework (which 50th anniversary occurs this year). This is achieved simply resetting the parameters of the variance reversion speed and the volatility of the variance. This reconciliation ultimately gives an idea of the potential pricing gap due to the choice of a different model and helps the trader to compare the prices of structured products valued according to different pricing approaches compared to the market standard (Giudici and Pagnottoni, 2019), (Giudici, Pagnottoni and Polinesi, 2020).

5) Conclusions

In this study the most popular numerical integration schemes for the Heston bivariate dynamical system (S_t, v_t) or $(\ln S_t, v_t)$ (Rouah, 2013) have been described and implemented. The methods which do not intrinsically admit the generation of negative values of the variance proved to be particularly interesting from a theoretical point of view. Among the analyzed methods, we cite the Transformed Volatility scheme (Zhu, 2010), the Balanced Implicit Scheme (Platen & Heath, 2009) and the Moment Matching (Andersen, 2008). Applied to the pricing of the most common investment certificates, the Moment Matching method has shown a lower convergence performance compared to the other techniques. Considering in addition the study by Kahl and Jäckel (2006) in which they showed that the convergence of the Balanced Implicit Scheme does not always prove to be uniformly valid, we can conclude that the Transformed Truncation Scheme has recorded results in line with theoretical expectations in the pricing of structured products and that it is also verified a priori that it is impossible for a negative variance to be generated in the asset projections. It is important to point out that there are other integration schemes for the Heston model in the literature, among which we mention the exact one proposed by Broadie and Kaya (2006). Indeed, the latter has a very important theoretical relevance as it ensures an exact numerical resolution of the integration problem, although it is difficult to implement in practice (Van Haastrecht and Pelsser, 2010) and the numerical processing times can be quite long (Bottasso *et al.*, 2023a).We consider it interesting for the continuation of the study to test the methods discussed in this paper, in particular the Transformed Volatility scheme and possibly others, that guarantee from a theoretical point of view the intrinsic positivity of the stochastic dynamics of the variance, on other types of certificates in order to generalize the conclusions in a multi-asset context.

References

[1] "Acepi associates' primary market 2023 Q1" Report from the main association specialized in Certificates.
[2] Albrecher H., Mayer P., Schoutens W., Tistaert J. (2007). "The Little Heston Trap". Wilmott Magazine, January 2007, 83-92.
[3] Andersen L. B. G. (2008). "Efficient Simulation of the Heston Stochastic Volatility Model." Journal of Computational Finance, 11(3):1–42.
[4] Andersen L.B.G., Brotherton-Ratcliffe R. (2005). ''Extended Libor Market Models with Stochastic Volatility.'' Journal of Computational Finance, 9(1):1–40.
[5] Begin J. F., Bedard M., Gaillardetz P. (2015). "Simulating from the Heston model: a gamma approximation scheme". Monte Carlo Methods and Applications Vol. 21, N. 3
[6] Bottasso A., Fusaro M., Giribone P. G., Tissone A. (2023a). "Implementation of variance reduction techniques applied to the pricing of investment certificates". Risk Management Magazine Vol.18, N. 1.
[7] Bottasso A., Fusaro M., Giribone P. G., Tissone A. (2023b). "Investment Certificates pricing using a Quasi Monte Carlo framework: case-studies based on the Italian market". International Journal of Financial Engineering. Forthcoming.
[8] Broadie M., Kaya O. (2006). "Exact Simulation of Stochastic Volatility and Other Affine Jump Diffusion Processes". Operations Research Vol. 54, N.2.
[9] Giribone P. G., Ligato S. (2013), "Metodologie per migliorare la velocità di convergenza nei simulatori Monte Carlo: Analisi delle tecniche ed implementazione in un framework di pricing". AIFIRM (Associazione Italiana Financial Industry Risk Managers) Magazine Vol. 8, N. 2.

[10] Giribone P. G., Ligato S. (2014), “Progettazione di un controllo affidabile sull'errore commesso dall'introduzione di sequenze a bassa discrepanza in un framework di pricing Monte Carlo” – AIFIRM (Associazione Italiana Financial Industry Risk Managers) Magazine Vol. 9, N. 1.
[11] Giribone P. G., Revetria R. (2021). “Certificate pricing using Discrete Event Simulations and System Dynamics theory”. Risk Management Magazine Vol. 16, N. 2.
[12] Giudici P., Pagnottoni P. (2019). “High frequency price change spillovers in bitcoin markets”. Risks, 7(4): 111
[13] Giudici P., Pagnottoni P., Polinesi G. (2020). “Network models to enhance automated cryptocurrency portfolio management”. Frontiers in Artificial Intelligence Vol. 3 p. 22
[14] Glasserman, P. (2003). Monte Carlo Methods in Financial Engineering. New York, NY: Springer.
[15] Heston S. L. (1993). ‘‘A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options.’’ Review of Financial Studies, 6:327–43.
[16] Kahl, C., Jäckel P. (2006). “Fast Strong Approximation Monte-Carlo Schemes for Stochastic Volatility Models.” Quantitative Finance, 6(6):513–36.
[17] Kloeden, P. E., Platen E. (1992). Numerical Solution of Stochastic Differential Equations. New York, NY: Springer.
[18] Lord R., Koekkoek R., van Dijk D. (2008), “A Comparison of Biased Simulation Schemes for Stochastic Volatility Models”. Tinbergen Institute Discussion Paper.
[19] Mrázek M., Pospíšil J. (2017). “Calibration and simulation of Heston model." Open Mathematics 15.1: 679-704.
[20] Platen, E., and D. Heath. (2009). A Benchmark Approach to Quantitative Finance, Volume 13. New York, NY: Springer.
[21] Rouah F. D. (2013). The Heston Model and its Extensions in Matlab and C#. First Edition. New Jersey. Wiley Finance.
[22] Van Haastrecht, A., Pelsser A. (2010). “Efficient, Almost Exact Simulation of the Heston Stochastic Volatility Model.” International Journal of Theoretical and Applied Finance, 13(1):1–43.
[23] Zhu, J. (2010). Applications of Fourier Transform to Smile Modeling: Theory and Implementation. Second Edition. New York, NY: Springer.