
Option Pricing in Practice—Heston’s Stochastic Volatility Model

Sascha Desmettre, Ralf Korn, and Tilman Sayer

1 The Finance Industry as Employer, Supplier of Mathematical Challenges, and Risk Factor

Over the past 50 years, the growth in complexity of the products offered in the financial markets has opened up a completely new sub-domain in mathematics—(modern) financial mathematics—which, in terms of its mathematical sophistication goes far beyond what used to be understood as financial mathematics, which would be better referred to as business accounting. The significant role of mathematics in the financial world can be demonstrated emphatically in a number of ways. The finance and insurance industries are today, more than ever, the most important employers of graduates from mathematics curricula, although thorough training in applied mathematics (especially stochastics, statistics, optimization, and numerics) is an advantage. In the financial field, mathematicians are sometimes employed as traders, as so-called *quants* (mathematics experts who implement quantitative methods and models in the investment banking sector), or as external staff for large consulting firms, who offer their competences to banks and insurance companies. The insurance branch even has its own professional designation—*actuary*—awarded by the Deutsche Aktuarvereinigung DAV after completion of an intensive training and examination.

The new mathematical challenges can be illustrated by means of four significant domains in financial mathematics:

S. Desmettre · R. Korn (✉) · T. Sayer
Abteilung Finanzmathematik, Fraunhofer ITWM, Fraunhofer Platz 1, 67663 Kaiserslautern,
Germany
e-mail: ralf.korn@itwm.fraunhofer.de

- *Modeling*—This is the simulation of price movements of all kinds (stock prices, interest rates, currency exchange rates, etc.). Here, stochastic processes are applied, based essentially on so-called *Itô processes*, developed in the middle of the last century. As functions of time, they are non-differentiable and require their own calculus, the *Itô calculus* (see [33] or [4] for an introduction).
- *Portfolio optimization*—The goal here is to determine an optimal trading strategy for an investor. Today, this represents one of the main fields of dynamic optimization and stochastic control (see [1] or [3]).
- *Risk management*—This deals with measuring and managing the risks associated with unforeseeable developments in investments. In the past years, the theory of risk measures has opened up a new field of theoretical research (see [15], the standard reference).
- *Option pricing*—This is the showpiece of financial mathematics. It deals with determining the prices of option contracts and has led to numerous new theoretical problems and results in such fields as martingale theory and the numerics of stochastic processes (see [4, 34], and [5]). In the following discussion, we will consider this field in greater depth.

Finally, due to the financial industry's intrinsic uncertainties, it represents one of the largest risk factors in modern economies, possibly *the* largest. The role of mathematics in this regard can indeed be viewed with some ambivalence. On the one hand, the methods of mathematical modeling and financial mathematics provide significant support in recognizing, understanding, and managing risks. On the other hand, the technology of mathematics can tempt its users into creating new, ever more complicated products, in the belief that mathematical methods can make their characteristics and risks understandable and controllable. Here, it is essential for mathematicians to emphasize the limits of modeling and, in particular, to point out the enormous dependency of the models on their input parameters—a risk that is often largely ignored.

Moreover, it is important to make clear the difference between prediction and simulation. Whereas stock price simulations can indeed be used to calculate risk measures or options prices by means of Monte Carlo methods, these simulations often have about as much utility in predicting actual stock prices as the choice of a number 2, . . . , 12 has in predicting the toss of a pair of dice. The financial crisis of 2007–2012 should be taken as a strong warning to sharpen the awareness of the limits of mathematical modeling and simulation. Conversely, however, one should also be aware that mathematical modeling was not the cause of the financial crisis. This fig leaf—so happily put to use by investment bankers—is not large enough to cover up the actual causes, such as egregious misjudgments of creditworthiness, the taking of tremendously risky market positions, and, in many cases, plain ignorance in acquiring products that were not understood.

In the course of this chapter, we will concentrate on the field of option pricing, which is a central concern of both modern investment banking and financial mathematics theory, but which is treated quite differently in these two fields. We want to demonstrate that the popular Black–Scholes model is no longer adequate for many practical purposes and

introduce instead the Heston model (see [29]), an alternative often used in practice that represents a compromise between practicability and theoretical generality. In so doing, we will encounter mathematical challenges in the areas of modeling, theoretical stochastics, and the numerical computation of options prices.

Along the way, we will first introduce the reader to the terminology of the world of options, describe project collaborations between the financial industry and the Financial Mathematics Department of the ITWM, and then delve into the theoretical foundations of stock price modeling and options pricing.

2 Options as Modern Ingredients of the Financial Markets

The terms *option* and *derivative* have developed negative connotations in the aftermath of the great financial crisis of 2007–2012. This is not entirely unjustified, given that many credit derivatives, in both their form and ultimate impact, were so complicated and opaque—not only for lay persons—that their trading contributed significantly to the outbreak of the crisis. Ironically, the designation *exotic option*, applied to many of these securities, seems very aptly chosen, given their mysterious and arcane nature.

In the following section, we first present the basic features of options and option trading by using the simplest options as examples. We then take a look at more complicated types, in order to illustrate the necessity of mathematical modeling.

2.1 Simple Options—Call and Put


The word *option*, in its colloquial sense, stands for an opportunity that one is not compelled to take, but which one may indeed take, if so desired. Having such an opportunity is always a good thing. Options contracts on the financial market, which represent a similar opportunity, are therefore to be had only at a price.

Option contracts are securities derived from underlying assets, which explains why they are also known as *derivatives*. These securities have been traded for centuries in one form or another, but they only achieved great economic significance at the beginning of the 1970s. They are used to secure market positions and design special payoff profiles, as well as for purposes of pure speculation. As the term itself implies, possession of an option includes a right of choice that the owner can, but does not have to, exercise.

With the simplest option, the *European call* on a stock, the buyer has the right (but not the obligation!) to acquire from the seller a share of a given company at time T (the *maturity*) at a fixed price K (the *exercise price* or *strike*). He will only do so if it is to his advantage, that is, if the stock price $S(T)$ lies above the strike K , since he could otherwise obtain the same share more cheaply by acquiring it directly in the market. Consequently, possession of a European call is equivalent to the payment

$$Y_{\text{call}} = (S(T) - K)^+ = \max(S(T) - K, 0)$$

Fig. 1 Final payoff of a European call with strike price K



at time T and is often depicted graphically by means of its payment profile, as in Fig. 1. Because the amount of the final payment is non-negative and may be positive, the possession of the option today must have a positive value. Determining this value is the object of option pricing.

The direct counterpart to the European call is the *European put*, which is given by the final payment

$$Y_{\text{put}} = (K - S(T))^+$$

at time T , and which gives the owner the right to sell a share at price K at time T to the seller of the European put. Here, the owner will only exercise his right if the current market price of the stock is lower than K . Such simple options as the European call and put are often referred to in the market as *vanilla options*.

L.F. Bachelier's dissertation *Théorie de la Spéculation* (see [16]), from 1900, may be regarded as the long-forgotten starting point for option pricing. Bachelier's idea was to use asset price modeling to derive theoretical values for different types of options on particular assets and to compare these values with the actual market prices. As the option price, he proposed using the expected value of the future payment arising from the option. In so doing, he used implicitly, for the first time, the so-called Brownian motion (with drift) as an asset price model (see Sect. 4.1)—albeit, without designating it as such.

One consequence of this modeling, however, was that the prices in the model could fall below zero. His ideas were taken up again in the 1960s with the introduction of geometric Brownian motion (see again Sect. 4.1) as a price model. In 1973, Fischer Black and Myron Scholes achieved the first crucial breakthrough in the field of option pricing with the derivation of explicit price formulas for European calls and puts (see [20] and Sect. 4.3).

2.2 Exotic Options—The Next Stage

For the simple options described above, the payment resulting from ownership depends only on the stock price $S(T)$ at the maturity of the option. This need not be the case, however. The general form of European options is given by a final payment of the form

$$Y = f(S(t), t \in [0, T]),$$

where f (and thus also Y) is a function describing the entire price development of the stock over the time interval $[0, T]$. Here, we identify time $t = 0$ as the present moment and $S_0 > 0$ as the initial stock price. One then speaks of a *path-dependent option*. For the European call and put, this dependency was given simply by the stock price at $t = T$. Examples of European options with genuine path dependency are:

- the *lookback* or *maximum option*, for which the maximum of the stock price appears in the final payment as follows:

$$Y_{lb} = \left(\max_{t \in [0, T]} S(t) - K \right)^+.$$

- the *Asian option*, which is given by a final payment in the form

$$Y_{ao} = \left(\frac{1}{T} \int_0^T S(t) dt - K \right)^+$$

for example, and similar variants (e.g., as discrete mean).

- or the *barrier option* class, for which the stock price in the interval $[0, T]$ may not exceed or must exceed (depending on the variant) one or more specified barriers H_i , so that at T , a positive payment flows to the option's owner. One example is the *double barrier knock out call*, with strike K , barriers $0 \leq H_1 < H_2$, and final payment


$$Y_{dbkoc} = 1_{\{H_1 < S(t) < H_2 \forall t \in [0, T]\}} (S(T) - K)^+.$$

Here, as seen in Fig. 2, there may indeed be price paths that end above the strike K , but violate a barrier condition beforehand and thus do not lead to a final payoff.

The pricing of such options that are strongly dependent on the stock price path—options that are often grouped together under the title of exotic options—requires a specialized and highly efficient numerical method and also demonstrates empirically the need for more sophisticated stock price models than the geometric Brownian motion (cf. here also Sect. 4.6).

Remark 1 It is certainly justified to ask why such complicated options are traded at all. The answer is multi-faceted.

Fig. 2 Two possible stock price paths and barriers $H_1 = 75$ and $H_2 = 170$, as well as strike $K = 100$ for a double barrier knock out call



A barrier option must obviously be cheaper than its counterpart without barriers since, in order to receive the same payment at the maturity of the option, the barrier conditions must be fulfilled, which is not always the case. If they are not fulfilled, the buyer of the barrier option receives nothing, whereas the buyer of the common European variants receives the full payment. Despite this risk, barrier options are happily bought in preference to the common variants, due to their lower price.

In the case of the maximum option, the option's owner wants to secure the largest possible difference between the option price and the execution price—an advantage for which he must then also be ready to pay a higher price than for the simple European call.

Finally, the Asian option represents a kind of insurance against short-term market manipulation. It is quite conceivable that large market players might use their trading power at the option's maturity to drive the price of the option's associated stock in a direction that works to their own advantage, but this tactic is not possible over a long time period. Therefore, Asian options typically use an average of stock prices over specified time points to determine the size of the option payment.

2.3 American Options and More—Free Choice

Another natural variant is the so-called *American option*, in which the option's holder can decide at what time-point $t \in [0, T]$ he exercises his option right. Here, the option can be identified by a whole family $Y(t)$, $t \in [0, T]$ of possible final payments, from which the option's owner can choose one. For each of the above-mentioned European option types, there is also a corresponding American variant, which, for a call, for example, is given by the following family of possible final payments:

$$Y_{\text{call}}(t) = (S(t) - K)^+, \quad t \in [0, T].$$

Generally speaking, there is a tremendous variety of options available in the financial marketplace. Along with options on stocks, one finds options on bonds, commodities, loans, options, contracts, foreign currencies, electricity, and virtually any asset that

is traded. Each of these option classes often generates its own mathematical problems for price calculation, not least because the underlying assets exhibit completely different characteristics. For example, electricity is, in general, non-storable, whereas commodities generate inventory costs but some may offer compensating strategic advantages, etc.

For those interested in more background on options trading, including historical background, we recommend [4] or [2].

2.4 How Much Do Options Cost?

This fundamental question regarding option pricing cannot be answered without two significant ingredients:

- a mathematical model of the price development of the asset underlying each option and
- the insight that, since an option is a derivative, its price must always depend on the current market price of the underlying asset.

Both ingredients are considered in detail in Sect. 4 and lead to a surprising result—one that was honored with the Nobel Prize for Economics.

3 Options at the ITWM

Because option pricing is one of the central domains of financial mathematics and so important for trading in modern financial markets, it also plays a central role in the work of the Financial Mathematics Department at the Fraunhofer ITWM. Our non-disclosure agreements with customers in the finance and insurance industries require us to remain somewhat vague in our descriptions of the associated projects, but we do want to give our readers an impression of what sort of work a mathematics research institute can undertake in this field.

The essential components underlying all the projects in the field of option pricing are:

- the development of new, and appropriate modification of existing, dynamic stock price models, toward the end of achieving realistic modeling (do the price movements in the model have the same characteristics as the empirically observed ones?) and numerical tractability (can parameters be stably and efficiently calculated?),
- the derivation of explicit analytical pricing formulas for special, exotic options,
- the development of numerical algorithms for pricing exotic options without explicit price formulas, for which new Monte Carlo methods, tree methods, and solution methods for differential equations are put to use,
- and the implementation of the developed algorithms in modern software for direct application by the trader or for risk management.

The years 2000 to 2013 witnessed the successful completion of many industrial projects. Among our clients were banks, such as the Hypovereinsbank and the Landesbank Baden-Württemberg; insurers, such as the R+V Versicherung; and financial services providers, such as Assenagon Asset Management S.A. The projects varied greatly in size and ranged from pricing a single options class to preparing complete software libraries for pricing exotic options. The latter was of a scale to incorporate all the above-mentioned research and development aspects and draw upon the complete spectrum of applied financial mathematics; it also led to publications in top-flight journals (see e.g., [6, 8, 26, 35], or [9]).

In order to remain continuously abreast of scientific developments in the field of option pricing, a multitude of PhDs were completed during this same time period relating to its various aspects. Often, algorithmic aspects were paramount, in order to make a satisfactory model fit for service in the first place. Consequently, in [11], the numerical pricing of so-called barrier options was investigated; in [12], new tree methods for pricing exotic options in the field of interest rates were developed; in [13], tree methods for option pricing in the Heston model were derived (see Sect. 6 also); and in [10], Monte Carlo methods for special multi-asset barrier options were examined.

4 The Foundations of Stock Price Modeling and Option Pricing

Modeling the development of the basic processes underlying the financial markets represents the foundation of financial mathematics. Depending on the market segment, these might be stock prices, indices (such as the DAX or Dow Jones), interest rates, exchange rates, or other indicators. In this section, we initially restrict ourselves to modeling only stock prices as stochastic processes and assume all other influencing quantities to be constant. For the technical bases, we refer the reader to [4].

4.1 Modeling Stock Prices

We wish to examine a financial market in which trading takes place in continuous time, any desired division of shares is permissible, and ancillary costs (broker and transaction fees, etc.) do not exist. As basis investment opportunities, we take the investment in a fixed-term deposit account and in (initially) one stock (or stock index).

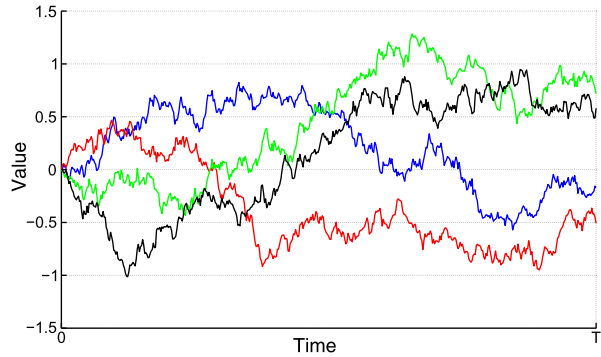
We assume that a fixed-term deposit $B(t)$ accrues interest continuously at a constant rate r , which leads to the temporal development

$$B(t) = B_0 e^{rt} \tag{1}$$

for an initial deposit of B_0 at initial time-point $t = 0$.

The significant ingredient for modeling the stock price is the selection of the stochastic process. Motivated by the central limit theorem, according to which the (centered and standardized) sum of many independent, identically-distributed random variables is asymptot-

Fig. 3 Four simulated paths of a Brownian motion with $n = 500$ and $T = 1$



ically normally distributed, we select a Brownian motion as the random driver of the stock price.

Definition 1 Let (Ω, \mathcal{F}, P) be a complete probability space. A Brownian motion $\{(W(t), \mathcal{F}_t), t \in [0, \infty)\}$ is a real-valued stochastic process with $W(0) = W_0 = 0$ and continuous paths with stationary and independent increments, that is

$$W(t) - W(s) \sim W(t - s) \quad \forall t > s \geq 0,$$

$$W(t) - W(s) \text{ is independent of } \mathcal{F}_u \text{ for } t \geq s \geq u \geq 0.$$

Here, $\{\mathcal{F}_t, t \in [0, \infty)\}$ is a right-continuous filtration for which \mathcal{F}_0 already contains all P -null sets. The filtration is called natural filtration if it is the filtration generated by the Brownian motion.

Remark 2 It can be shown that the requirements placed on the Brownian motion in Definition 1 already yield $W(t) \sim \mathcal{N}(0, t)$. From this, an algorithm follows directly to (approximately) simulate a Brownian motion. To this end, for $n \in \mathbb{N}$ and $T > 0$, let $0 = t_0 < t_1 < \dots < t_n = T$ define a separation of the interval $[0, T]$. We then proceed as follows:

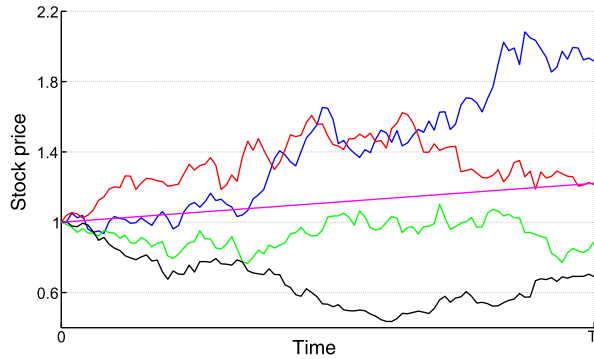
1. Set $W(0) = 0$.
2. Generate n independent $\mathcal{N}(0, 1)$ -distributed random numbers Z_1, \dots, Z_n .
3. For $i = 1, \dots, n$, set

$$W(t_i) = W(t_{i-1}) + \sqrt{t_i - t_{i-1}} Z_i$$

and interpolate linearly between $W(t_i)$ and $W(t_{i-1})$.

Several of the (discretized) paths generated for $T = 1$, $n = 500$, and $t_j = j/n$ according to the above algorithm are shown in Fig. 3. Here, another characteristic of the paths of Brownian motion can also be detected; they are nowhere differentiable with respect to time. This is very significant for modeling a stock price as a function $S_t = f(W_t)$, since

Fig. 4 Simulated stock price paths and mean $\mathbf{E}(S(t))$ for $b = 0.2$ and $\sigma = 0.4$



this is then also not differentiable with respect to time. This characteristic is indispensable from a modeling perspective, if one sticks to the continuity of the price over time. Were the stock price to be differentiable with respect to time, then it would also be locally predictable, and, as a result, no trading would take place.

With the help of the Brownian motion, the stock price $S(t)$ is modeled as a *geometric Brownian motion*

$$S(t) = S_0 e^{(b - \frac{1}{2}\sigma^2)t + \sigma W(t)}, \tag{2}$$

that is, its logarithmized increments are assumed to be normally distributed. Here, b and σ are real numbers that describe the *mean rate of return* and the *volatility* of the stock price. Furthermore,

$$\begin{aligned} \mathbf{E}(S(t)) &= S_0 e^{bt}, \\ \mathbf{Var}(S(t)) &= S_0^2 e^{2bt} (e^{\sigma^2 t} - 1), \\ \ln\left(\frac{S(t)}{S_0}\right) &\sim \mathcal{N}\left(\left(b - \frac{1}{2}\sigma^2\right)t, \sigma^2 t\right). \end{aligned}$$

Figure 4 shows the price paths $S(t)$ associated with the paths of the Brownian motion for the parameters $b = 0.2$, $\sigma = 0.4$, and $S_0 = 1$, along with the course of the mean $\mathbf{E}(S(t))$.

Remark 3 To formulate the model thus derived for multiple stocks, one introduces, for d stocks, an n -dimensional ($n \geq d$) Brownian motion $W(t) := (W^{(1)}(t), \dots, W^{(n)}(t))$, whose components are each independent, one-dimensional Brownian motions according to Definition 1, and models the price of the j -th stock $S^{(j)}(t)$ as

$$S^{(j)}(t) = S_0^{(j)} e^{(b^{(j)} - \frac{1}{2} \sum_{k=1}^n \sigma_{j,k}^2)t + \sum_{k=1}^n \sigma_{j,k} W^{(k)}(t)}, \quad j = 1, \dots, d,$$

where $b^{(j)}$, $\sigma_{j,k}$, and $S_0^{(j)}$ are suitable constants. Due to the characteristics of the normal distribution, the stock prices remain log-normally distributed, and the expectations and variances can also be determined analogously.

As our next ingredient, we introduce the investors by means of the trading strategy, where the information structure of the investors is given by the filtration $\{\mathcal{F}_t\}_{t \in [0, T]}$ corresponding to the Brownian motion. Here, a trading strategy is a two-dimensional stochastic process, whose components specify the number of units of each security being held.

Definition 2

(a) A *trading strategy* φ is an \mathbf{R}^2 -valued process $\varphi(t) := (\varphi_0(t), \varphi_1(t))'$ that is progressively measurable with regard to $\{\mathcal{F}_t\}_{t \in [0, T]}$. Moreover, we require

$$\int_0^T |\varphi_0(t)| dt < \infty \quad P\text{-almost surely,}$$

$$\int_0^T (\varphi_1(t)S(t))^2 dt < \infty \quad P\text{-almost surely.}$$

The value $x := \varphi_0(0)B_0 + \varphi_1(0)S_0$ is called initial value of φ .

(b) Let φ be a trading strategy with initial value $x \geq 0$. The process

$$X(t) := \varphi_0(t)B(t) + \varphi_1(t)S(t)$$

is then called the *wealth process* corresponding to φ with *initial wealth* $X(0) = x$.

(c) A trading strategy φ is called *self-financing* if, for the associated wealth process $X(t)$, $t \in [0, T]$,

$$X(t) = x + \int_0^t \varphi_0(s)dB(s) + \int_0^t \varphi_1(s)dS(s)$$

P -almost surely, that is, the current wealth is yielded by the sum of the initial wealth and the profits/losses from investments in the time period $[0, t]$. It is then called *admissible* when its associated wealth process is non-negative.

Note that the requirement of progressive measurability of the strategy means that the investor has no information about the future development of the stock price. The economically natural requirement that the investor behaves in a self-financing way is a genuine requirement and does not result mathematically from parts (a) and (b) of the definition. For information on the analogous definition of a trading strategy in more generalized markets (see above), we refer the reader to [4].

4.2 Option Pricing and the Arbitrage Principle

With the mathematical market model developed in the previous section, we are now in a position to tackle the problem of option pricing. The essential idea behind option pricing is, first, that an option is a derived security (derivative) having no existence of its own independent from its underlying asset; the movement of the asset price also determines

the price of the option. The second central principle that applies in option pricing is that of *absence of arbitrage*. Here, one considers an arbitrage opportunity to be a transaction involving the possibility of a profit without the risk of a loss, where none of the investor's own money must be used.

A typical example of such an arbitrage opportunity is a free ticket in a lottery. Here, although one may seldom win, neither must one put up one's own money to acquire the ticket.

An arbitrage opportunity of this type is such a good "deal" that every market participant would instantly take advantage of it. The resulting infinite demand would immediately trigger a corresponding price adjustment on the market and the arbitrage opportunity would disappear. Therefore, for theoretical deliberations, one only considers financial market models that are free from arbitrage opportunities. This assumption alone makes it possible in any arbitrary financial market model to set lower and upper bounds for option prices (as a function of each option type). See, for example, Chap. 3 in [4].

Definition 3 An arbitrage opportunity is an admissible trading strategy φ whose associated wealth process $X(t)$ fulfills the conditions

$$X(0) = 0, \quad X(T) \geq 0 \quad P\text{-almost surely}, \quad P(X(T) > 0) > 0.$$

In the market we are considering here, with prices according to Eqs. (1) and (2), an even stronger variant of the absence-of-arbitrage principle can be shown, namely, the validity of the *replication principle*. This stipulates that two investments with identical future cash flows must have the identical price today. If this were not so, one could buy the cheaper of the two alternatives today and simultaneously sell the more expensive. The future cash flows arising from this transaction neutralize one another, but one has already accrued today an increase in wealth from the price difference, which could then simply be invested in a money market account. Because one needs no starting capital to pursue this strategy, but the final wealth is strictly positive, this represents an arbitrage opportunity.

The following is a central theorem in the theory of option pricing.

Theorem 1 (Completeness of the market) *Using the notation $\theta = (b - r)/\sigma$, then $H(t) := \exp(-(r + \theta^2/2)t - \theta W(t))$.*

(a) *Let $x \geq 0$. For an admissible trading strategy φ with wealth process $X(t)$, we have*

$$\mathbf{E}(H(t)X(t)) \leq x \quad \forall t \in [0, T].$$

(b) *Let $Y \geq 0$ be a \mathcal{F}_T -measurable random variable with*

$$\tilde{x} := \mathbf{E}(H(T)Y) < \infty. \tag{3}$$

Then, there is an admissible trading strategy φ with initial value \tilde{x} , wealth process $X(t)$ and

$$X(T) = Y \quad P\text{-almost surely.}$$

The complete market theorem seems unspectacular at first blush, but it is extremely significant for option pricing. Part (b) says that a non-negative final payment Y (that fulfills condition (3)), which can be secured via possession of an option, for example, can be synthetically generated by pursuing a suitable trading strategy φ in the money market account and the stock. Thus, from the perspective of the final payment, it is irrelevant whether one physically possesses the option or whether it is synthetically replicated. If there is no arbitrage opportunity in the market, then Eq. (3) yields the option price \tilde{x} . The non-existence of an arbitrage opportunity follows immediately from part (a) of the theorem, however (see [4]).

Corollary 1 (Absence of arbitrage) *In the market model under consideration here, there is no arbitrage opportunity.*

Consequently, Theorem 1 and Corollary 1 together lead directly to the main result for option pricing in this market model.

Corollary 2 (Fair price) *In the market model under consideration here, the fair price of an option with final payment Y , which is compatible with the arbitrage principle, is given by*

$$x_Y := \mathbf{E}(H(T)Y) \quad (4)$$

when this value is finite.

Remark 4 (Option price, risk-neutral pricing, equivalent martingale measure) If one considers $H(t)$ to be a discount factor process, with which one discounts future payments, then Eq. (4) says that one obtains the price of the option with final payment Y by calculating the expectation of the final payment discounted by $H(T)$. This means, first, that the net present value principle, under which the price is defined as the future payment discounted to today, is valid here with a stochastic discount factor. It can also be shown, however, that

$$\mathbf{E}(H(T)Y) = \mathbf{E}_Q(\exp(-rT)Y) \quad (5)$$

holds true, where the second expectation with regard to the (unique) probability measure Q is formed in the probability space being considered, for which

$$\mathbf{E}_Q(\exp(-rt)S(t)) = S_0$$

holds true. Because, as a consequence, $S(t)/S_0$ and $B(t)/B_0$ possess the same expectation under Q , regardless of whether one is dealing with a risky or a risk-free investment, Q is also called the *risk-neutral measure*, and one speaks of *risk-neutral pricing*, since the

option price is given by the right side of Eq. (5). The existence of the risk-neutral measure Q follows from the Girsanov theorem (see [4], Chap. 3), which also states that, under Q , the process

$$W_Q(t) = W(t) + \theta t,$$

with θ from Theorem 1, is a Brownian motion. If one inserts this in the stock price equation (2), then one obtains

$$S(t) = S_0 \exp\left(\left(r - \frac{1}{2}\sigma^2\right)t + \sigma W_Q(t)\right),$$

from which it once again follows, with the defining characteristics of the Brownian motion, that the discounted stock price process $S(t)/B(t)$ is a martingale. Since it is also true that the measure Q is equivalent to P (i.e., both measures possess the same null sets), one also refers to Q as the *equivalent martingale measure*. The relationship between the existence of such equivalent martingale measures and the absence of arbitrage in a market model is also referred to as the first fundamental theorem of option pricing (see [25]). In general, it can be shown in an elementary fashion for analogous, arbitrage-free financial market models that establishing an option price by the right side of Eq. (5) does not lead to arbitrage opportunities when Q is an equivalent martingale measure.

4.3 The Black–Scholes Formula: Nobel Prize for Mathematics

For the special case of the European call option, one can explicitly calculate the expectation that determines the option price. This then yields the famous *Black–Scholes formula* (see [20] or [4]).

Theorem 2 (Black–Scholes formula) *In the market model given by the price equations (1) and (2), the price $X_{\text{call}}(t, S(t), K, T)$ of a European call option at time $t \in [0, T]$ with maturity T and strike $K > 0$ is given by*

$$X_{\text{call}}(t, S(t), K, T) = S(t)\Phi(d_1(t)) - Ke^{-r(T-t)}\Phi(d_2(t)) \quad (6)$$

where $\Phi(\cdot)$ is the distribution function of the standard normal distribution and where we use the abbreviations

$$d_1(t) = \frac{\ln(S(t)/K) + (r + \sigma^2/2)(T - t)}{\sigma\sqrt{T - t}}, \quad d_2(t) = d_1(t) - \sigma\sqrt{T - t}.$$

Using the same notation, the price of the corresponding European put is given by

$$X_{\text{put}}(t, S(t), K, T) = Ke^{-r(T-t)}\Phi(-d_2(t)) - S(t)\Phi(-d_1(t)).$$

Remark 5 (Characteristics, applications, and consequences of the Black–Scholes formula)

- (a) The outstanding quality of the Black–Scholes formula is not simply that it allows a closed analytical form of the price of the European call, but that it allows this price to be independent from b , the stock's mean rate of return. Because this parameter is far more critical for estimating the future development of the stock price, which can only be poorly estimated from past stock prices (although one can efficiently estimate the volatility σ , at least as σ^2 , from historical data), it is precisely the absence of b that is one of the main reasons for the market's acceptance of the Black–Scholes formula—aside from its elegant and convincing mathematical derivation. Its standing was further underscored in 1997 with the awarding of the Nobel Prize for Economics to Robert C. Merton and Myron Scholes for their work on it. Fischer Black had already died in 1995 and could therefore no longer be honored for his contribution.
- (b) In the market, the Black–Scholes formula is not generally used to calculate call prices, but rather in a manner indicating that the market does not fully believe in the Black–Scholes model. Closed pricing formulas, including those in other models, are frequently used for *parameter calibration*, that is, the input parameters for each model are defined so that the associated model prices for each derivative coincide as well as possible with the prices observed in the market. In the case of the Black–Scholes formula, this is taken a step further, in that a positive volatility is defined for all calls with different maturities and different strikes on the same stock, such that the price observed in the market coincides exactly with the model price (see Sect. 4.3.1). This volatility is called the *implied volatility* of the particular call. If one joins the resulting points by means of a suitable interpolation procedure, one then obtains a so-called *implied volatility surface*. If the Black–Scholes model corresponded exactly to the market data, then all of the implied volatilities would have to be (at least) virtually identical. In the following section, we make clear that this is not so by defining in detail the implied volatility and implied volatility surface and illustrating them with an example.


Other weaknesses of the Black–Scholes model regarding characteristic empirical properties of stock and option prices (so-called *stylized facts*) are treated in Sect. 4.6.

4.3.1 Implied Volatility

According to Theorem 2, the price of a European call is a function with six arguments, all of which are observable except for the volatility. Consequently, if the volatility is known, the option price can be calculated. Conversely, if the option price is known, one can easily show that the volatility can be uniquely determined under the assumption that it is positive. Therefore, it is possible to determine the implied volatility σ_{imp} from the option prices quoted in the market.

For $i = 1, \dots, N$, we let $X_{\text{call}}^{\text{market}}(K_i, T_i)$ denote the market price of a European call with exercise price K_i and maturity T_i , where the same strike may very well be paired with different maturities and vice versa. If one sets these market prices equal to the theoretical

Fig. 5 Implied volatility surface on 14 December 2011, from European calls on the stock of Allianz SE



prices of the corresponding call options in the Black–Scholes model, then σ_{imp} can be uniquely determined from

$$X_{\text{call}}(t, S(t), K_i, T_i) \stackrel{!}{=} X_{\text{call}}^{\text{market}}(K_i, T_i),$$

for $i = 1, \dots, N$. Because the market prices are dependent on the exercise prices and maturities, the implied volatilities are dependent on them also.

Definition 4 The representation of the implied volatility σ_{imp} as a function of the exercise price K and the maturity T is referred to as the *implied volatility surface*.

Figure 5 shows the implied volatility surface on 14 December 2011, as obtained from European calls on the stock of Allianz SE. As the graphic shows, contrary to the assumption in the Black–Scholes model, options having different execution prices and maturities possess different implied volatilities.

4.4 Alternative Stock Price Models: Theoretical Aspects

There are several ways to redress the deficits of the Black–Scholes model, and these are often resorted to when modeling problems. Among others, they are:

- Moving from a linear to nonlinear stochastic dynamics, as introduced, for example, in Sect. 4.4.1.
- Introducing further stochastic components, such as a stochastic, rather than constant, volatility (see Sect. 4.4.2 and, particularly, Sects. 5 and 6).
- Considering a more general class of stochastic processes for modeling the uncertainty in the stock price process, such as the class of Lévy models in Sect. 4.4.3.

4.4.1 Local Volatility Models

Local volatility models utilize the first of the above-mentioned ideas. To avoid the problem of non-constant volatility, the volatility of the stock price is permitted to be time and location dependent. As before, a simple one-dimensional Brownian motion $W(t)$ is used as the underlying stochastic process. This is done in the hope of thereby maintaining the completeness of the market. The replication principle of option pricing would then remain valid. In point of fact, an astounding result is attained in this regard, which we will present in Theorem 3.

We consider a market model consisting of the usual money market account with interest rate r (see Eq. (1)) and a stock whose price is modeled with the aid of the stochastic differential equation

$$dS(t) = rS(t)dt + \sigma(S(t), t)S(t)dW(t), \quad S(0) = S_0. \quad (7)$$

Here, we let $\sigma(x, t)$ be a non-negative, real-valued function of such a form that Eq. (7) possesses a unique (non-negative) solution. One sees immediately that, for the constant function $\sigma(x, t) \equiv \sigma$, one obtains the Black–Scholes model.

Now, instead of prescribing a parametric form of the volatility function, Dupire [27] takes an entirely different approach. Motivated by the terminology of the implied volatility surface, he looks for a volatility function that ensures, for a specified set of call prices, that the associated theoretical option prices (calculated as the discounted expectation of the final payoff under the unique equivalent martingale measure) coincide with the given market prices. And this is precisely the assertion of the following theorem.

Theorem 3 ([27]) *Let today's market prices $X_{\text{call}}^{\text{market}}(0, S, K, T)$ of European calls for all possible choices of strikes $K \geq 0$ and maturities $T \geq 0$ be known, be once differentiable as functions of the maturity, and be twice differentiable as functions of the strike. With the choice of the volatility function $\sigma(x, t)$ via*

$$\sigma(K, T) = \frac{1}{K} \sqrt{\frac{2 \frac{\partial X_{\text{call}}^{\text{market}}}{\partial T} + rK \frac{\partial X_{\text{call}}^{\text{market}}}{\partial K}}{K^2 \frac{\partial^2 X_{\text{call}}^{\text{market}}}{\partial K^2}}}, \quad (8)$$

the market prices coincide with the theoretical call prices obtained in the corresponding local volatility model according to

$$X_{\text{call}}(0, S, K, T) = \mathbf{E}(e^{-rT}(S(T) - K)^+) \quad \forall (T, K) \in [0, \infty)^2.$$

Here, it is implicitly assumed that the call prices are furnished in such a way that all expressions appearing in Eq. (8) are defined.

Theorem 3 presents exactly the desired result. Consequently, there exists for any given set of market prices for European calls, a volatility function $\sigma(x, t)$ that generates them. Thus, one has found a model in which the theoretical model prices coincide with the given

market prices for simple options. It is therefore plausible to use this model for calculating the prices of more complicated options for which there are no market prices. The problem with the theorem, however, is its practical applicability; some prerequisites and assumptions that enter the result cannot be verified and/or can hardly be implemented in practice:

- To design the volatility function, one needs a continuous set of market prices. Due to the discreteness of the set of strikes and maturities, however, there is none. Therefore, the volatility function must be obtained with the help of interpolation and extrapolation methods, but is then dependent on the method being used and, in particular, is no longer unique.
- In a local volatility model generated in this fashion, there are generally no closed, analytical price formulas, even for simple standard options.
- The form of the local volatility function has no intuitive economical interpretation or motivation, but is based purely on data.

For further general aspects, we refer the reader to [27].

A popular parametrical model, which, however, represents no substantial improvement over the Black–Scholes model, is the CEV model (*Constant-Elasticity-of-Variance* model), for which the stock price equation is given as

$$dS(t) = rS(t)dt + \sigma S(t)^\alpha dW(t), \quad S(0) = S_0$$

with $\alpha \in [0, 1]$ and $r, \sigma \in \mathbf{R}$. For the special choice of $\alpha = 0$ and $\alpha = 1$, it admits explicit solutions:

- For $\alpha = 1$, one then obtains the already familiar geometric Brownian motion (Black–Scholes case), that is, log-normally distributed stock prices.
- For $\alpha = 0$, one obtains

$$S(t) = S_0 \exp(rt) + \sigma \int_0^t \exp(r(t-u)) dW(u),$$

from which follows that the stock price is normally distributed with

$$\mathbf{E}(S(t)) = S_0 \exp(rt), \quad \mathbf{Var}(S(t)) = \frac{\sigma^2}{2r} (\exp(2rt) - 1).$$

For all values $\alpha \in [0, 1)$, the CEV model admits a quite complicated, albeit closed, formula for the price of a European call (see [5]), which we will not reproduce here. The additional parameter α does indeed yield, in comparison with the Black–Scholes model, a somewhat better fit to option market prices, but one that is still far from perfect. Moreover, for $0 < \alpha < 1$, the model is numerically difficult to manage. For these reasons, we do not recommend it for practical application.

4.4.2 Stochastic Volatility Models

The economic idea behind stochastic volatility models is that price fluctuations are determined by supply and demand and, depending on the trading intensity, may be stronger or weaker. Since the intensity of the price fluctuations in the Black–Scholes model is determined by the value of the constant volatility σ , one assumes here a trading intensity that is (on average) constant.

If, on the other hand, one wishes to model a non-constant trading intensity whose variability cannot be predicted, then it makes sense to model the volatility by a stochastic process also. Such a *stochastic volatility model* is then given by price and variance process equations having the form

$$dS(t) = bS(t)dt + \sqrt{v(t)}S(t)dW_1(t), \quad S(0) = S_0, \quad (9)$$

$$dv(t) = \alpha(t)dt + \beta(t)dW_2(t), \quad v(0) = v_0, \quad (10)$$

where $\alpha(t)$ and $\beta(t)$ are suitable stochastic or deterministic processes that are progressively measurable relative to the filtration generated from the two-dimensional Brownian motion $(W_1(t), W_2(t))$. Furthermore, we let v_0 be the initial value of the variance process and $\rho \in [-1, 1]$ be the correlation of the Brownian motions $W_1(t)$ and $W_2(t)$,

$$\text{Corr}(W_1(t), W_2(t)) = \rho. \quad (11)$$

Analogously to σ in the Black–Scholes model, we call $\sqrt{v(t)}$ the *volatility process*. Moreover, all of the processes described above should be selected so that the coupled stochastic differential equations (9) and (10) possess a unique solution.

In practice, one tends to be less interested in the economic motivation behind stochastic volatility models. The decisive factors are the free parameters and/or processes arising from the introduction of the stochastic differential equation (10), with whose help one hopes to obtain a model that can much more accurately replicate the option prices observed in the market.

Among the various choices found in the literature for modeling the volatility process, the choice of Heston (see [29]) has proved especially effective in practice and has, in many fields, replaced the Black–Scholes model as the standard. At the ITWM, we have already successfully applied the Heston model in several industrial projects. In Sects. 5 and 6, we offer an extensive theoretical description of the model and take a closer look at the details of its application for modeling variants and pricing algorithms.

4.4.3 Lévy Models

In the class of Lévy models, a *Lévy process* $Z(t)$ essentially takes over the role of the Brownian motion $W(t)$ from the Black–Scholes model. A Lévy process is a stochastic process with independent and stationary increments that starts with $Z(0) = Z_0 = 0$ and possesses paths that are almost surely continuous. Thus, a Brownian motion is also a Lévy process, but a significant majority of Lévy processes possess paths exhibiting jumps. Lévy models

are determined by their characteristics. They typically exhibit a large number of parameters and their distributions, in comparison with a normal distribution, possess markedly sharper densities with heavier tails. These can therefore explain even extreme stock price movements, for which the Black–Scholes model has no explanation (or only an explanation such as: “In the credit crisis, we have observed 10σ -events”).

For an overview of the application of Lévy processes in financial mathematics, we refer the reader to the monographs [23] and [40]. Other models known in the theory that have also been applied to market data include the hyperbolic model (see [28]), the variance gamma model (see [37]), and the NIG model (see [17]). To date, however, the Lévy models have been unable to make large-scale breakthroughs in practical application, since the extensive parameterization is connected with greater estimation effort and larger estimation error.

4.5 Further Application Aspects


The given application is a crucial factor in choosing a stock price model. A simple model, such as the Black–Scholes, often suffices to price relatively simple derivatives. For complicated, strongly path-dependent exotic options, however, the Black–Scholes model is generally inadequate. It is somewhat paradoxical, then, that when pricing options based on multiple assets (so-called *basket options*), one often resorts to the Black–Scholes model again in its multi-dimensional variant. The explanation here is that there are no suitable multi-dimensional variants of the above-mentioned, more realistic models, or none that would be numerically and statistically manageable.

Finally, the computation time required to determine an individual option price is another crucial argument. Banks often carry out *sensitivity analyses* when selling large amounts of a particular derivative. This involves varying all possible input parameters, which can quickly lead to an exponentially increasing number of different scenarios, for which the option prices must then be calculated. Hence, research into faster algorithms and new hardware concepts, such as the use of graphic cards or so-called FPGA as computational accelerators, remains an active field.

4.6 Effects with Real Data: Stylized Facts as an Argument Against the Black–Scholes Model

Stock prices, interest rates, exchange rates, and many other financial time series exhibit typical empirical characteristics that distinguish them from other time series. These characteristics are referred to as stylized facts. In the following analysis, we will present those characteristics in particular which suggest that the assumption of constant volatility in the Black–Scholes model is too restrictive. Here, we take the term *discrete time series* to mean an ordered sequence of observations at discrete time-points, such as exists with stock prices, for example.

Fig. 6 Daily log returns of the DAX between January 2008 and December 2013



Definition 5 Let $S(t)$ be the price of a stock. We define the return $R(s, t)$ between the time points s and $t > s$ as

$$R(s, t) := \frac{S(t) - S(s)}{S(s)},$$

and the logarithmized return (log return) $r(s, t)$ as

$$r(s, t) := \ln\left(\frac{S(t)}{S(s)}\right).$$

With regard to a stock's daily return, we define $r(n) := r(n - 1, n)$ for $n \in \mathbf{N}$.

Remark 6 In the following discussion, we present the typical characteristics of the daily, and thus discrete, time series of the log return $r(n)$, $n \in \mathbf{N}$. For small price changes, as are the norm with stock data, the log returns are a good approximation of the returns. The time series relevant to the investigation relate to the daily closing prices between January 2008 and December 2013.

Let the sample mean $\hat{\mu}_N$ and the sample variance $\widehat{\mathbf{Var}}_N$ of the log return be defined as

$$\hat{\mu}_N = \frac{1}{N} \sum_{n=1}^N r(n), \quad \widehat{\mathbf{Var}}_N = \frac{1}{N-1} \sum_{n=1}^N (r(n) - \hat{\mu}_N)^2.$$

4.6.1 Volatility Clustering

Figure 6 shows the daily log returns of the DAX for the relevant time period. The graphic clearly illustrates that there are phases with both large and small price changes, which alternate with each other. This phenomenon is referred to as volatility clustering.

4.6.2 The Leverage Effect

Empirical data shows that, for returns on stocks, negative reports in the form of higher losses have a stronger impact on the perception of risk (and thus of volatility) than positive reports in the form of higher profits. The volatility thus reacts asymmetrically to the

Table 1 Sample skewness $\hat{\gamma}_{N=1527}$ of the DAX and various DAX stocks for log returns from January 2008 through December 2013

Sample skewness	
DAX	0.1028
BASF	0.0594
BMW	0.0861
Deutsche Bank	0.3028
Deutsche Telekom	-0.0379

signs of shocks. This phenomenon is known as the leverage effect. In 1976, Fischer Black commented on this as follows: “A drop in the value of the firm will cause a negative return on its stock, and will usually increase the leverage of the stock. [...] That rise in the debt-equity ratio will surely mean a rise in the volatility of the stock.” Therefore, price and volatility changes are usually negatively correlated.

4.6.3 The Skewness—A Measure for the Symmetry of a Distribution

The empirical distribution of logarithmized stock price returns is often asymmetric. One measure for this asymmetry is the skewness of a random variable.

Definition 6 Let X be a real-valued random variable with $\mathbf{E}(X^3) < \infty$. The *skewness* $\gamma(X)$ of X is defined as

$$\gamma(X) := \frac{\mathbf{E}((X - \mathbf{E}(X))^3)}{(\mathbf{Var}(X))^{3/2}}.$$

Remark 7 For the discrete log returns, the skewness is estimated by means of the sample skewness

$$\hat{\gamma}_N = \frac{1}{\widehat{\mathbf{Var}}_N^{3/2}} \frac{1}{N} \sum_{n=1}^N (r(n) - \hat{\mu}_N)^3.$$

The sample skewness of a normally distributed random variable is equal to zero. The more $\hat{\gamma}_N$ deviates from zero, the more asymmetric is the empirical distribution of the data. If $\hat{\gamma}_N < 0$ (left-skewed), the left tail of the distribution is heavier than the right. Conversely, for $\hat{\gamma}_N > 0$ (right-skewed), the right tail is heavier than the left.

Table 1 shows the sample skewness of the DAX and some of its individual components. All observed values are non-zero and the associated time series are accordingly asymmetric. This in turn suggests considering alternative stock price models that do not assume a normal distribution.

4.6.4 Kurtosis—Emphasized Peaks and Tails

Figure 7 shows the histogram of the log returns of the DAX for the relevant time period, along with the density of the adjusted normal distribution. As the graphic indicates, the density of the log returns has a higher peak in the middle and heavier tails than the density

Fig. 7 Empirical distribution of the DAX log returns and density of the fitted normal distribution





Fig. 8 Q-Q plot of the log returns for the DAX



of the normal distribution. The quantile-quantile diagram (Q-Q plot) in Fig. 8 makes clear how heavy the tails of the empirical distribution are in comparison to the normal distribution. If the historical data had been normally distributed, it would lie on the dashed red line.

Definition 7 Let X be a real-valued random variable with $\mathbf{E}(X^4) < \infty$. The *kurtosis* $\kappa(X)$ of X is defined as

$$\kappa(X) := \frac{\mathbf{E}((X - \mathbf{E}(X))^4)}{\mathbf{Var}(X)^2}.$$

Remark 8 The kurtosis for the discrete log returns is estimated on the basis of the sample kurtosis

$$\hat{\kappa}_N = \frac{1}{\widehat{\mathbf{Var}}^2} \frac{1}{N} \sum_{n=1}^N (r(n) - \hat{\mu}_N)^4.$$

Normally distributed random variables have a kurtosis of 3. If the kurtosis is larger, then the distribution of the associated random variable is leptokurtic. The distribution then has a narrower peak than that of a normal distribution.

Table 2 Sample kurtosis $\hat{\kappa}_{N=1527}$ for the DAX and various DAX stocks for the log returns from January 2008 through December 2013

Sample kurtosis	
DAX	8.6948
BASF	10.1601
BMW	6.5074
Deutsche Bank	9.3832
Deutsche Telekom	12.8838

Table 2 shows the sample kurtosis of the DAX and various stocks for the time period January 2008 through December 2013. All observed values are significantly larger than 3; the associated time series thus exhibit pronounced tails and high peaks. These characteristics are typical for mixtures of distributions with different variances. Therefore, these results also indicate that the assumption of constant volatility is not appropriate.

4.6.5 The Volatility Reverts to Its Mean

Another empirical characteristic of the volatility is that it reverts to its mean. To investigate this behavior, we consider the historical standard deviation of the log returns. This is referred to as the historical volatility.

Definition 8 The *historical N -days volatility* σ_{hist} is defined as the annualized standard deviation

$$\sigma_{\text{hist}}(N) := \sqrt{\frac{D}{N-1} \sum_{n=1}^N (r(n) - \hat{\mu}_N)^2}.$$

Here, D stands in general for a days convention, which specifies the number of days used to approximate a year, since weekends and holidays cause the exact number to fluctuate. In practice, $D = 252$ is often used.

In order to study the historical volatility, we consider the rolling historical volatility over a longer time period.

Definition 9 For $l \in \mathbf{Z}$, one takes the *rolling historical N -days volatility* to be the time series

$$\sigma_{\text{hist}}(N, l) := \sqrt{\frac{D}{N-1} \sum_{n=l+1}^{l+N} (r(n) - \hat{\mu}_N(l))^2},$$

where the sample mean $\hat{\mu}_N(l)$ on the basis of N is calculated for the observed data points, starting at 1, and then slides over the data.

Fig. 9 Historical rolling one-year volatility for the DAX

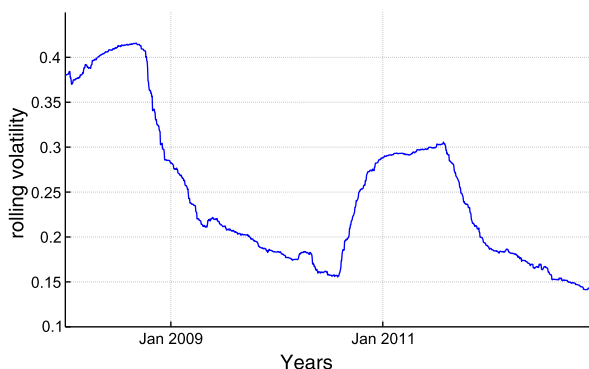


Figure 9 shows the rolling historical volatility on a one-year basis $\sigma_{\text{hist}}(252, l)$ for the DAX from January 2008 through January 2013. One can observe that the historical volatility, after reaching high (low) values, tends to fall (climb). Empirically, the volatility reverts to its mean.

In summary, one can state that both the stylized facts and the implied volatility observed in the market (see Sect. 4.3.1) militate against the assumption of constant volatility in the model. Instead, the volatility itself should be modeled as a random variable that is correlated with the stock price. One model that does so is Heston’s stochastic volatility model (cf. [29]); this will be analyzed in depth in the following discussion, along with its variants—some of which we have put to use in industrial projects.

5 Theoretical Foundations of the Heston Model

The Heston model is a stochastic volatility model in which the functions $\alpha(t)$ and $\beta(t)$ from Eq. (10) possess a special form. Here, the stock price and the variance both follow the stochastic differential equations

$$dS(t) = bS(t)dt + \sqrt{\nu(t)}S(t)dW_1(t), \quad S(0) = S_0, \quad (12)$$

$$d\nu(t) = \kappa[\theta - \nu(t)]dt + \sigma\sqrt{\nu(t)}dW_2(t), \quad \nu(0) = \nu_0. \quad (13)$$

As in Eq. (11), the Brownian motions $W_1(t)$ and $W_2(t)$ have a correlation of ρ . Moreover, b denotes the stock drift; κ , the reversion speed of the variance to the mean reversion level $\theta > 0$; and σ , the volatility of the variance. The process $\nu(t)$ from Eq. (13) is called the square root diffusion process, or Cox-Ingersoll-Ross (CIR) process. It is the path-wise unique, weak solution of Eq. (13) and is almost surely non-negative. It is not given explicitly, but has a non-central chi-square distribution and, in particular, is finite. If the Feller condition

$$2\kappa\theta \geq \sigma^2 \quad (14)$$

also holds, then the process is strictly positive, that is, $P(v(t) > 0) = 1 = Q(v(t) > 0)$ for all $t \geq 0$. Furthermore, the variance process reverts to its mean reversion level θ , which—as described in Sect. 4.6—is an empirical characteristic of the volatility. The correlation of the Brownian motions is in a position to replicate the leverage effect described earlier, and is thus generally negative (sometimes even very close to -1 !). All in all, the Heston model thus models all the characteristics of the volatility that were described as stylized facts.

As does $W_1(t)$, the Brownian motion $W_2(t)$ also represents a source of uncertainty. However, because the volatility is not an asset that can be traded in the market, the replication principle—which is based on the completeness of the market (see Theorem 1)—can no longer be applied. In such an incomplete market, the risk-neutral pricing measure Q is no longer unique. Moreover, there are infinitely many equivalent martingale measures (see [19] or [18]).

Up to this point, the Heston model has been considered under the physical measure P , which is supposed to describe the price movements in the real market. The dynamics under an equivalent martingale measure Q can be derived from the dynamics (12) and (13). For a positive constant λ , the risk-neutral parameters

$$\kappa^* = \kappa + \lambda, \quad \theta^* = \frac{\kappa\theta}{\kappa + \lambda},$$

and the Girsanov transformations

$$\begin{aligned} dW_1^Q(t) &= dW_1(t) + (b - r) \int_0^t \frac{1}{\sqrt{v(s)}} ds, \\ dW_2^Q(t) &= dW_2(t) + \frac{\lambda}{\sigma} \int_0^t \sqrt{v(s)} ds, \end{aligned}$$

can be used to define the risk-neutral form of the Heston model as follows:

$$dS(t) = rS(t)dt + \sqrt{v(t)}S(t)dW_1^Q(t), \quad S(0) = S_0, \tag{15}$$

$$\begin{aligned} dv(t) &= [\kappa(\theta - v(t)) - \lambda v(t)]dt + \sigma\sqrt{v(t)}dW_2^Q(t) \\ &= \kappa^*[\theta^* - v(t)]dt + \sigma\sqrt{v(t)}dW_2^Q(t), \quad v(0) = v_0. \end{aligned} \tag{16}$$

Here, $W_1^Q(t)$ and $W_2^Q(t)$ denote Q -Brownian motions with correlation ρ .

Remark 9 In Heston’s original work (cf. [29]), the term $\lambda v(t)$ is referred to as the *market price of the volatility risk* Φ . This (and therefore the associated Girsanov transformation, also) can be *a priori* freely selected. Both economic and mathematical arguments militate for modeling proportional to variance $v(t)$; only for this choice is there a known semi-closed formula for the price of European calls and puts.

In closing, we want to point out that the choice of a particular equivalent martingale measure equates to the choice of a market price for the volatility risk, which is ultimately determined by the choice of the positive constant λ . Consequently, we must also pose the question of which measure is to be used in the specific application. The answer to this question is revealed in Sect. 6.1.

5.1 Closed Form Solution for the Price of European Calls

One of the main reasons for the success of the Heston model in practice is a semi-closed price formula for European calls and puts that allows one to efficiently determine the model parameters from market prices, and thus, to calibrate the model (see Sect. 6.1). Using classical arbitrage arguments, one obtains the following partial differential equation for determining the price of a European call $X_{\text{call}}(t, S, K, T)$:

$$0 = \frac{\partial X_{\text{call}}}{\partial t} + \frac{\nu S^2}{2} \frac{\partial^2 X_{\text{call}}}{\partial S^2} + \rho \sigma \nu S \frac{\partial^2 X_{\text{call}}}{\partial \nu \partial S} + \frac{\sigma^2 \nu}{2} \frac{\partial^2 X_{\text{call}}}{\partial \nu^2} + r S \frac{\partial X_{\text{call}}}{\partial S} - r X_{\text{call}} + [\kappa(\theta - \nu) - \lambda \nu] \frac{\partial X_{\text{call}}}{\partial \nu}, \quad (17)$$

where it is assumed that the market price of the volatility risk is proportional to the variance, according to the relationship $\Phi = \lambda \nu(t)$. There is no known explicit solution for the partial differential equation (17). However, Heston found a way to express the solution with the aid of characteristic functions. Analogously to the Black–Scholes formula (6), he chooses the approach

$$X_{\text{call}}(t, S(t), K, T) = S(t) P_1(S(t), \nu(t), t, \ln(K)) - K e^{-r(T-t)} P_2(S(t), \nu(t), t, \ln(K))$$

for the solution, where $P_1(S(t), \nu(t), t, \ln(K))$ and $P_2(S(t), \nu(t), t, \ln(K))$ describe the probabilities that the stock finishes above the strike. Both probabilities fulfill the partial differential equation. If the characteristic functions $\varphi_1(S(t), \nu(t), t, u)$ and $\varphi_2(S(t), \nu(t), t, u)$ belonging to the probabilities exist, then $P_1(S(t), \nu(t), t, \ln(K))$ and $P_2(S(t), \nu(t), t, \ln(K))$ are given by their inverse Fourier transforms

$$P_j(S(t), \nu(t), t, \ln(K)) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \Re \left[\frac{e^{-iu \ln(K)} \varphi_j(S(t), \nu(t), t, u)}{iu} \right] du \quad (18)$$

for $j = 1, 2$, where $\Re(\cdot)$ denotes the real part. The linearity of the coefficients then suggests the approach

$$\varphi_j(S(t), \nu(t), t, u) = \exp(C_j(\tau, u) + \nu D_j(\tau, u) + iu \ln(S(t))), \quad (19)$$

for $j = 1, 2$ and $\tau := T - t$ for the characteristic functions. Utilizing $\varphi_1(S(t), v(t), t, u)$ and $\varphi_2(S(t), v(t), t, u)$ in Eq. (17) then delivers the following system of linear differential equations

$$0 = -\frac{u^2}{2} + \rho\sigma ui D_j + \frac{\sigma^2}{2} D_j^2 + u_j ui - b_j D_j - \frac{\partial D_j}{\partial \tau}, \quad (20)$$

$$0 = rui + aD_j - \frac{\partial C_j}{\partial \tau} \quad (21)$$

for the unknowns $C_j(\tau, u)$ and $D_j(\tau, u)$ with initial conditions

$$C_j(0, u) = 0, \quad D_j(0, u) = 0 \quad (22)$$

and

$$u_1 = \frac{1}{2}, \quad u_2 = -\frac{1}{2}, \quad a = \kappa\theta, \quad b_1 = \kappa + \lambda - \rho\sigma, \quad b_2 = \kappa + \lambda. \quad (23)$$

The solution of the system (20), (21) and (22) is given by

$$C_j(\tau, u) = rui\tau + \frac{a}{\sigma^2} \left[(b_j - \rho\sigma ui + d_j)\tau - 2 \ln \left[\frac{1 - g_j e^{d_j \tau}}{1 - g_j} \right] \right],$$

$$D_j(\tau, u) = \frac{b_j - \rho\sigma ui + d_j}{\sigma^2} \left[\frac{1 - e^{d_j \tau}}{1 - g_j e^{d_j \tau}} \right] \quad (24)$$

with

$$g_j = \frac{b_j - \rho\sigma ui + d_j}{b_j - \rho\sigma ui - d_j}, \quad d_j = \sqrt{(\rho\sigma ui - b_j)^2 - \sigma^2(2u_j ui - u^2)}. \quad (25)$$

The following theorem summarizes the results.

Theorem 4 (Heston's price formula) *Let the market price of the volatility risk be given by $\Phi = \lambda v(t)$. Then, in the Heston model, which is specified by Eqs. (12), (13), and (11), the arbitrage-free price of a European call is given by*

$$X_{\text{call}}(t, S(t), K, T) = S(t)P_1(S(t), v(t), t, \ln(K)) - Ke^{-r(T-t)}P_2(S(t), v(t), t, \ln(K)).$$

The probabilities $P_j(S(t), v(t), t, \ln(k))$ and the associated characteristic functions $\varphi_j(S(t), v(t), t, u)$ are given by Eqs. (18) and (19). The further quantities are defined in Eqs. (23), (24), and (25).

5.2 Variants of the Heston Model—Requirements Arising from Practice

On the basis of the acceptance and popularity of the Heston model in practice, the Financial Mathematics Department of the Fraunhofer ITWM received numerous research commissions from the financial and insurance industries, whose goals were the model's theoretical generalization and algorithmic implementation. In the wake of these projects, new and innovative variants of the closed formula from Theorem 4 were developed and implemented. In this section, we treat several of these variants—particularly those that resulted in publications in relevant journals.

5.2.1 The Heston Model with Time-Dependent Coefficients

The partial differential equation (20) is a nonlinear differential equation of the Riccati type. Therefore, generalizing the Heston model for non-constant parameters is non-trivial. The work of Mikhailov and Nögel (cf. [38]) considers diverse variants for treating time-dependent coefficients. For example, since Eq. (20) is not dependent on the mean reversion level θ , a general solution for a time-dependent enhancement $\theta(t)$ can be found. Other special cases include solutions with the help of hyper-geometric functions, for cases in which the reversion speed is modeled as $\kappa(t) = at + b$ or $\kappa(t) = ae^{-at}$. Strictly speaking, however, one must resort to other techniques. By numerically solving Eqs. (20) and (21), the model's application can be extended with relative ease to the situation of time-dependent parameters. Here, Runge–Kutta algorithms are good candidates. The use of semi-closed price formulas arises for the algorithmic implementation—especially for calibrating the model.

Asymptotic Expansion Because an analytical solution for the partial differential equation (20) can only be found for a few special cases, it seems appropriate to apply asymptotic methods. We therefore assume that $\rho(t)$ results from a superposition of time-dependent functions and, for small variations ϵ , possesses a potential series expansion around a constant value ρ_0 :

$$\rho(t) = \rho_0 + \epsilon\rho_1(t) + \epsilon^2\rho_2(t) + \dots$$

Using the approach

$$D_j(t) = D_{j,0}(t) + \epsilon D_{j,1}(t) + \epsilon^2 D_{j,2}(t) + \dots$$

the first order approximation delivers a linear equation with time-dependent coefficients, whose solution is given by

$$D_{j,1}(t) = -\sigma u_j i \int_0^t \rho_1(\tau) D_{j,0}(\tau) \exp\left(\int_0^\tau D_{j,0}(\xi) d\xi - (-\rho_0 \sigma u_j i + b_j) \tau\right) d\tau \\ \times \exp\left(-\int_0^t D_{j,0}(\tau) d\tau + (-\rho_0 \sigma u_j i + b_j) t\right).$$

As an alternative to the above asymptotic approach, one could perform an asymptotic analysis of the system with slowly changing parameters.

Piece-Wise Constant Parameter If one sub-divides the time interval $[t, T]$ into n sub-intervals $[t, t_1], \dots, [t_i, t_j], \dots, [t_{n-1}, T]$ and defines the model parameters to be constant in each sub-interval, then a closed solution can be found for Eq. (20), even for different parameters in different sub-intervals. With the help of the time inversion $\tau_k = T - t_{n-k}$, $k = 1, \dots, n - 1$, the initial condition for the first sub-interval $[0, \tau_1]$ is exactly zero. For this interval, one can then use the solution (24) of the Heston model. For the second sub-interval, we need solutions for the differential equations (20) and (21) with arbitrary initial conditions

$$C_j(0, u) = C_j^0, \quad D_j(0, u) = D_j^0, \tag{26}$$

which are given by

$$\begin{aligned} C_j(\tau, u) &= rui\tau + \frac{a}{\sigma^2} \left[(b_j - \rho\sigma ui + d_j)\tau - 2 \ln \left[\frac{1 - g_j e^{d_j \tau}}{1 - g_j} \right] \right], \\ D_j(\tau, u) &= \frac{b_j - \rho\sigma ui + d_j - (b_j - \rho\sigma ui + d_j)g_j e^{d_j \tau}}{\sigma^2(1 - g_j e^{d_j \tau})} \end{aligned} \tag{27}$$

with

$$g_j = \frac{b_j - \rho\sigma ui + d_j - D_j^0 \sigma^2}{b_j - \rho\sigma ui - d_j - D_j^0 \sigma^2}, \quad d_j = \sqrt{(\rho\sigma ui - b_j)^2 - \sigma^2(2u_j ui - u^2)} \tag{28}$$

and (23). The continuity requirement for the functions $C_j(\tau, u)$ and $D_j(\tau, u)$ at the intersection of the first and the second sub-interval τ_1 delivers the initial conditions for the second sub-interval as

$$C_j(0, u) = C_j^0 = C_j^H(\tau_1, u), \quad D_j(0, u) = D_j^0 = D_j^H(\tau_1, u), \tag{29}$$

where $C_j^H(\tau_1, u)$ and $D_j^H(\tau_1, u)$ refer to the Heston solution with the initial conditions (22). If one solves the above equations relative to the initial conditions C_j^0 and D_j^0 , one obtains the initial conditions for the second sub-interval. The procedure is then repeated for each jump point of the parameters τ_k , for $k = 2, \dots, n - 1$. Summarizing, the calculation of the option price in the Heston model with piece-wise constant parameters consists of 2 phases:

1. Determine the initial conditions for each sub-interval with the aid of the formulas in (29).
2. Determine the functions $C_j(\tau, u)$ and $D_j(\tau, u)$ using the solutions (27) and (28) with the initial conditions (26).

5.2.2 Forward Starting Options in the Heston Model

For pricing many exotic options in the Heston model, one must often resort to numerical methods, such as a Monte Carlo simulation or tree method (cf. Sect. 6). There are also instances, however, where closed formulas have been derived for complex derivatives. One example is the so-called *forward starting option*, which is treated in the work of Kruse and Nögel (cf. [35]).

A forward starting option is one whose exercise price is not completely determined until a time-point t^* . This time-point lies between the issuing date and the option's maturity, and is referred to as the starting point. Here, one can see that the forward starting option belongs to the class of path-dependent options. The payoff function for this option is given by

$$Y_{\text{fso}} = (S(T) - kS(t^*))^+, \quad (30)$$

where $k \in [0, 1]$ denotes a percentage.

Using the principle of risk-neutral pricing, a semi-closed pricing formula can be obtained for the option. The derivation goes beyond this discussion, however, so that we refer the interested reader to [35] for more information and present only the result here.

Theorem 5 (Forward starting option in the Heston model) *Let $\kappa \geq \rho\sigma$ and $0 \leq t < t^* < T$. If the stock price and the variance fulfill the risk-neutral dynamics (15) and (16), and if the Feller condition (14) also holds, then the price of a forward starting option at time t with payoff (30) is given by*

$$X_{\text{fso}}(t, S(t), K, T) = S(t)\hat{P}_1(t) - ke^{-r(T-t^*)}S(t)\hat{P}_2(t), \quad (31)$$

where

$$\hat{P}_j(t) := \int_0^\infty P_j(1, \xi, t^*, k) p(\xi, v(t)) d\xi$$

and the probabilities P_j are given in Eq. (18). Moreover,

$$p(\xi, v(t)) = \frac{B}{2} e^{-(B\xi + \Lambda)/2} \left(\frac{B\xi}{\Lambda} \right)^{(R/2-1)/2} I_{R/2-1}(\sqrt{\Lambda B\xi}) 1_{\{\xi > 0\}},$$

$$\Lambda = B e^{-(\kappa - \rho\sigma)(t^* - t)} v(t), \quad (32)$$

$$B = \frac{4(\kappa - \rho\sigma)}{\sigma^2} (1 - e^{-(\kappa - \rho\sigma)(t^* - t)})^{-1}, \quad (33)$$

and

$$R = \frac{4\kappa\theta}{\sigma^2},$$

where $I_{R/2-1}(\cdot)$ denotes the modified Bessel function of the first kind.

For a forward starting option on the return of a stock with payoff function

$$Y_{\text{rfsso}} = \left(\frac{S(T)}{S(t^*)} - K \right)^+, \tag{34}$$

a corresponding variant of the option price formula (31) can be specified on the basis of

$$\left(\frac{S(T)}{S(t^*)} - K \right)^+ = \frac{(S(T) - KS(t^*))^+}{S(t^*)}.$$

In the Heston model, the option price belonging to the payoff (34) is given by

$$X_{\text{rfsso}}(t, S(t), K, T) = e^{-r(t^*-t)} (\hat{P}_1(t) - Ke^{-r(T-t^*)} \hat{P}_2(t)), \tag{35}$$

where the expression $\kappa - \rho\sigma$ is replaced by κ in Eqs. (32) and (33).

Remark 10 For the numerical implementation of the option price formulas (31) and (35), we refer in particular to the calculation of the modified Bessel function of the first kind, which can be approximated by the following series expansion:

$$I_{R/2-1}(\sqrt{\Lambda B \xi(t^*)}) \approx \sum_{n=0}^N \frac{(\Lambda B \xi(t^*))^n}{2^{2n} n! \Gamma(n + R/2)}.$$

For practical applications, it turns out that the series converges with sufficient speed, so that even relatively small values of N are acceptable.

With the aid of the closed formulas (31) and (35), we have the efficient tools we need in order to price forward starting options.

5.2.3 A Sparsely Parameterized Multi-Asset Heston Model

In order to price options based on several underlying assets, a multi-dimensional version of the Heston model was developed at the Fraunhofer ITWM by Dimitroff, Lorenz, and Szimayer (cf. [26]). We now wish to present this work. We first point out that, with the help of the Cholesky decomposition, the risk-neutral dynamics of the Heston model (15) and (16) can be represented as follows:

$$\begin{aligned} dS(t) &= rS(t)dt + \sqrt{\nu(t)}S(t)dW(t), & S(t) &= S_0, \\ dv(t) &= \kappa[\theta - \nu(t)]dt + \sigma\sqrt{\nu(t)}[\rho dW(t) + \sqrt{1 - \rho^2} d\tilde{W}(t)], & \nu(t) &= \nu_0, \end{aligned}$$

where, for simplicity's sake, we dispense with the notation \star and Q , and start directly with the risk-neutral parameterization relevant for the pricing.

Multi-Dimensional Generalization In the following treatment, we describe a parsimonious, multi-dimensional extension of the one-dimensional Heston model, in which each one-dimensional sub-model is a classical one-dimensional Heston model, although the

price processes may exhibit correlations to each other. Consequently, the model is parsimonious in the sense that for a d -dimensional model, only $d(d - 1)/2$ correlations between the risky securities are needed. For $i = 1, \dots, d$,

$$\begin{aligned} \begin{pmatrix} dS_i(t) \\ dv_i(t) \end{pmatrix} &= \begin{pmatrix} r S_i(t) \\ \kappa_i(\theta_i - v_i(t)) \end{pmatrix} dt \\ &+ \begin{pmatrix} \sqrt{v_i(t)} S_i(t) & 0 \\ 0 & \sigma_i \sqrt{v_i(t)} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \rho_i & \sqrt{1 - \rho_i^2} \end{pmatrix} \begin{pmatrix} dW_i(t) \\ d\tilde{W}_i(t) \end{pmatrix} \end{aligned} \quad (36)$$

denotes the Heston model in vectorized form, where $W_i(t)$ and $\tilde{W}_i(t)$ describe the uncorrelated Brownian motions.

The model is thus defined, except for its dependency structure. Let $W(t) = (W_1(t), \dots, W_d(t))$ and $\tilde{W}(t) = (\tilde{W}_1(t), \dots, \tilde{W}_d(t))$ now be d -dimensional Brownian motions. For $i = 1, \dots, d$ and $j = 1, \dots, d$, we assume that $W(t)$ and $\tilde{W}(t)$ are described by the following dependency structure:

1. $W(t)$ has the correlation matrix $\Sigma^S = (\rho_{i,j})$, i.e., $\langle W_i(t), W_j(t) \rangle = \rho_{i,j}$,
2. $\tilde{W}(t)$ has the correlation matrix I_d , i.e., $\langle \tilde{W}_i(t), \tilde{W}_j(t) \rangle = \delta_{i,j}$,
3. $W(t)$ and $\tilde{W}(t)$ are independent.

The complete correlation matrix of $(W(t), \tilde{W}(t))$ is thus given by

$$\Sigma = \Sigma^{(W, \tilde{W})} = \begin{pmatrix} \Sigma^S & 0 \\ 0 & I_d \end{pmatrix}. \quad (37)$$

The first assumption allows for an arbitrary correlation structure between the risky securities. In contrast, the second and third assumptions stipulate that the dependency structure of the variance processes is determined by the corresponding correlations of the Brownian motions, which are transferred to the variance processes by the parameters ρ_i and ρ_j .

The model specification (36) and the assumed form of the correlation matrix (37) thus define the following correlation structure:

$$\begin{aligned} \frac{dS_i(t)dS_j(t)}{\sqrt{(dS_i(t))^2(dS_j(t))^2}} &= \rho_{i,j}, \\ \frac{dS_i(t)dv_j(t)}{\sqrt{(dS_i(t))^2(dv_j(t))^2}} &= \rho_{i,j}\rho_j, \\ \frac{dv_i(t)dv_j(t)}{\sqrt{(dv_i(t))^2(dv_j(t))^2}} &= \begin{cases} \rho_{i,j}\rho_i\rho_j, & \text{for } i \neq j, \\ 1, & \text{for } i = j. \end{cases} \end{aligned}$$

Remark 11 The one-dimensional models presented here $(S_i(t), v_i(t))$ are affine with the corresponding closed formulas, according to Theorem 4. However, the multi-dimensional generalization is not affine, and as a consequence, its characteristic function cannot be simply determined. Therefore, Monte Carlo methods and tree methods—as described in Sect. 6—are generally required for pricing options with multiple underlying investment assets.

Empirical Correlations and Correlation Adjustment Under the assumption that the parameters of the one-dimensional sub-models are known, there are additional $(d - 1)d/2$ free parameters from the matrix Σ^S that must be determined in order to correlate the risky securities. If there is sufficient data available, this is accomplished with the help of the implied correlations of multi-asset options. If this data is not available, the empirical correlations $\widehat{\Sigma}^{\text{emp}}$ from the time series of the risky securities can be estimated and adjusted to the model correlations Σ^S . Here, it is known that $\widehat{\Sigma}^{\text{emp}}$ is an unbiased estimator for the correlation matrix Σ^{emp} of the investment assets, which is evidently strongly dependent on the non-observed quantity Σ^S .

The idea is now to adjust the correlation matrix Σ^S so that it fits Σ^{emp} , which, in turn, is estimated by $\widehat{\Sigma}^{\text{emp}}$. Here, it is important to point out that Σ^S describes the infinitesimal correlation of the Brownian motion $W(t)$ and Σ^{emp} describes the correlation of the log returns. We refer to the adjustment of Σ^S to Σ^{emp} as the *correlation adjustment*. In the following treatment, we now formally define the estimator $\widehat{\Sigma}^{\text{emp}}$. Let $r_i(k)$ for $k = 1, \dots, K$ be discrete log returns of the i -th stock. Moreover, let

$$\widehat{v}_{i,j,T,K}^{\text{emp}}(\Sigma) = \frac{1}{K-1} \sum_{k=1}^K (r_i(k) - \widehat{\mu}_K^i)(r_j(k) - \widehat{\mu}_K^j). \tag{38}$$

Then, the empirical correlation matrix of the log returns is defined as

$$\widehat{\Sigma}_{T,K}^{\text{emp}}(\Sigma) = (\widehat{\rho}_{i,j,T,K}^{\text{emp}}(\Sigma))_{1 \leq i,j \leq d}$$

and its elements, as

$$\widehat{\rho}_{i,j,T,K}^{\text{emp}}(\Sigma) = \frac{\widehat{v}_{i,j,T,K}^{\text{emp}}(\Sigma)}{\sqrt{\widehat{v}_{i,i,T,K}^{\text{emp}}(\Sigma)\widehat{v}_{j,j,T,K}^{\text{emp}}(\Sigma)}}. \tag{39}$$

It can now be shown that the entries $\widehat{\rho}_{i,j,T,K}^{\text{emp}}(\Sigma)$ of the empirical correlation matrix converge suitably to the entries $\rho_{i,j}$ of the model correlation matrix Σ^S . That is, the model correlations Σ can be determined by calculating the historical, empirical correlations $\widehat{\Sigma}_{T,K}^{\text{emp}}(\Sigma)$ using (38) and (39). This makes it possible to develop a procedure for estimating the unknown correlations. If we assume that the empirical correlations are observed under the risk-neutral measure Q , then, if T and K are large, the observed sample correlations are good approximations for the expected sample correlations, given the true correlation

structure of the Brownian motions; that is,

$$\widehat{\Sigma}_{T,K}^{\text{emp}} \approx \mathbf{E}^Q \widehat{\Sigma}_{T,K}^{\text{emp}}(\Sigma^{\text{true}}) =: \Sigma^Q(\Sigma^{\text{true}}).$$

The unknown correlations can thus be determined by means of a minimization problem:

$$\min_{\Sigma \in \text{Cor}(d)} \left\| \Sigma^Q(\Sigma) - \widehat{\Sigma}_{T,K}^{\text{emp}} \right\|, \quad (40)$$

where $\text{Cor}(d)$ denotes the space of the $d \times d$ -dimensional correlation matrices and $\|\cdot\|$, a suitable matrix norm. The solution of the minimization problem (40) is not trivial; however, it can be solved using standard software. We denote the solution as Σ^* .

Generating an Admissible Correlation Matrix It is possible that the correlations estimated with the above algorithm may not lead to a valid (positive semi-definite) correlation matrix. In this case, a transformation is required. One possible algorithm that generates a genuine correlation matrix from an estimated one is the following (see [32] also):

1. Determine an eigenvalue decomposition of Σ^* as $\Sigma^* = S\Lambda S^T$, where $\Lambda = \text{diag}(\lambda_i)$.
2. Define the diagonal matrix $\tilde{\Lambda}$ with entries

$$\tilde{\lambda}_i = \begin{cases} \lambda_i & \text{if } \lambda_i \geq 0 \\ 0 & \text{if } \lambda_i < 0. \end{cases}$$

3. Generate the diagonal matrix T with entries

$$t_i := \left(\sum_m s_{im}^2 \tilde{\lambda}_m \right)^{-1}.$$

4. Define $B := \sqrt{T}S\sqrt{\tilde{\Lambda}}$ and obtain a new positive semi-definite correlation matrix as $\hat{\Sigma}^* := BB^T$ with $\hat{\Sigma}_{ii}^* = 1$.

For other relevant algorithms, we refer to [39], for example. Finally, then, with the generation of the correlation matrix, the sparsely parameterized multi-asset Heston model is completely defined.

6 The Heston Model in Action—Algorithmic Implementation

In this section, we turn to the questions that are relevant for implementing the Heston model.

6.1 Problems of Calibration

As previously shown, in a complete, arbitrage-free market, a derivative can be uniquely replicated by other investments available in the market. Therefore, in the theory of financial mathematics, the equivalent martingale measure and/or the market price of risk, is uniquely given by the model. As a consequence, the price of the derivative is also uniquely determined.

Because Heston's stochastic volatility model defines an incomplete financial market, the absence of arbitrage alone here does not suffice to uniquely determine a price; there are infinitely many equivalent martingale measures that define infinitely many arbitrage-free product prices. So-called lower and upper arbitrage bounds can then be specified for a financial product and, ultimately, all prices lying within these bounds are correct—according to financial mathematics theory.

In practice, these price bounds are insufficient. For the specific pricing of products, a single equivalent martingale measure must be chosen, which raises the following interesting question:

“Who determines the martingale measure?”

The short and amazing answer is (cf. [19]):

“The market does!”

The implication of the answer is simple: in determining the measure, one should include information available in the market in the form of traded products. This process, known as *model calibration*, uses the option prices observed in the market as input parameters. The goal is to use them to determine the model parameters so that the model prices correspond as closely as possible to the observed market prices.

However, since the number of traded products typically exceeds the number of model parameters by a wide margin, it frequently happens that not all market prices can be replicated exactly. The following algorithm uses the least squares method to calibrate the model for European calls.

For $i = 1, \dots, N$, let $X_{\text{call}}^{\text{market}}(K_i, T_i)$ be the prices of N European calls observed in the market for various exercise prices K_i and maturities T_i , and let $\omega_1, \dots, \omega_N$ be positive weights that add up to 1. We then obtain the simple *calibration algorithm* for the parameters $(v_0, \kappa, \theta, \sigma, \rho)$ describing the Heston model:

1. Solve the minimization problem

$$\min_{(v_0, \kappa, \theta, \sigma, \rho)} \sum_{i=1}^N \omega_i \left(X_{\text{call}}^{\text{market}}(K_i, T_i) - X_{\text{call}}(t, S(t), K_i, T_i) \right)^2.$$

The (calibrated) parameter set found here offers the best possible explanation for the observed market situation.

Here, one sees the decisive advantage of the Heston model. Since there are semi-closed calculations for the prices of European options, the required model prices do not have to be determined by means of laborious methods. In each iteration of the minimization algorithm, the N model prices can thus be obtained very quickly.

Remark 12

- (a) Because the above minimization problem is highly nonlinear, one needs methods of nonlinear optimization to find a solution. Here, one must take particular care that the solution algorithm for the global optimization problem can terminate in a local minimum. This makes it absolutely essential to check the resulting parameters for plausibility and, if needed, to start the optimization again using different initial values or different minimization algorithms.
- (b) There are both deterministic and stochastic algorithms available for solving the optimization problem, and each type has specific advantages and disadvantages. For example, deterministic algorithms lend themselves to situations in which good initial values for the calibration exist. Based on the initial solution, these then attempt to minimize the target function by locally changing the parameters. As a result, the deterministic methods often converge very quickly, but do not leave the neighborhood of a local optimum. In contrast here, the stochastic optimization methods offer the possibility of abandoning an already discovered local minimum and continuing the search for a better solution. Implementing these algorithms is typically more laborious, but the calibration results are often superior to those obtained via deterministic methods.
- (c) In addition to the option prices observed in the market, the market prices of other products can be used for calibration purposes. If these products do not have closed form solutions in the model, however, laborious numerical simulations are needed to determine the prices, and these are frequently very time-consuming. Therefore, the market prices of derivatives for which analytical solutions exist in the model form the basis for a satisfactory model calibration.
- (d) For practically relevant applications, the prices observed in the market exert differing influences on the calibration. This might be a function of the product-specific bid/ask spread, for example, which is a sign of a product's liquidity. For this reason, when calibrating, practitioners often use various weights ω_i to weight the individual input prices, in order to emphasize relevant situations or reduce the influence of less significant ones.

For realistic applications, the calibrated parameters typically vary over time. This means that it may be necessary to re-calibrate the model repeatedly within a short time period (within a single day, for example). For these applications, the calibrated parameters are often used as the new initial values for the re-calibration.

6.2 Pricing Complicated Products; Aspects of Numerical Simulation

In practice, pricing simple products such as European calls and puts is generally not a problem. In the following section, we consider the pricing of more complicated derivatives using numerical methods such as Monte Carlo simulations and tree approximations. While the Monte Carlo simulation for determining option prices is based on the strong law of large numbers, tree pricing relies on the central limit theorem. Each method has its advantages and disadvantages, so that, in practice, it has proved to be effective to implement both methods. In addition to these methods, there are other numerical methods, such as those for solving partial differential equations or Fourier techniques. We will not discuss these further here, but instead, refer the interested reader to [4] or [22] for more information.

6.2.1 Variants of the Euler Discretization

In order to price complex products traded in non-liquid markets, it is necessary to simulate the stock and variance paths of the Heston model.

Although the variance does not have a closed solution, its distribution—the non-central chi-squared distribution—is known. Thus, a promising approach might be to exactly simulate variance values $v(t)$ directly with the aid of the distribution. Such an approach is presented in [21]. With the exactly simulated variance process, the stock price process can then be determined using a suitable discretization method.

These methods function well for independent and therefore uncorrelated Brownian motions. However, problems arise in the generalized case for high absolute values of the correlation. For this case, an unbiased method is described in [21] that includes an inverted Fourier transformation. However, this method is much more time intensive than simpler discretization methods (see [36]).

The following algorithm introduces a naive discretization method suited to the Heston model that is based on the Euler–Maruyama method for the numerical solution of stochastic differential equations.

1. Initialize the variance and stock price approximation by $v(0) = v_0$ and $S(0) = S_0$.
2. Define $\Delta = T/n$, where T denotes the product maturity and n the number of discretization steps.
3. Repeat for $j = 1, \dots, n$:
 - (a) Simulate independent random variables $Z_1, Z_2 \sim \mathcal{N}(0, 1)$.
 - (b) Define $Z_3 = \rho Z_1 + \sqrt{1 - \rho^2} Z_2$.
 - (c) Discretize the stochastic differential equation of the variance and iterate

$$v(j\Delta) = v((j-1)\Delta) + \kappa(\theta - v((j-1)\Delta))\Delta + \sigma\sqrt{v((j-1)\Delta)}\Delta Z_3.$$

- (d) Discretize the stochastic differential equation for the logarithmized stock price $\mathcal{X}(t) = \ln(S(t))$ and iterate

$$\mathcal{X}(j\Delta) = \mathcal{X}((j-1)\Delta) + \left(r - \frac{v((j-1)\Delta)}{2}\right)\Delta + \sqrt{v((j-1)\Delta)\Delta}Z_1.$$

4. Determine the path $\mathcal{X}(t)$ as a linear approximation between the discrete time-points $\mathcal{X}(j\Delta)$ for $j = 0, 1, \dots, n$. Then, $S(t) = \exp(\mathcal{X}(t))$ is the stock price path.

Although the continuous-time solution of the variance process assumes only non-negative values, the approximation can indeed generate negative values. Here, however, the root terms that must be determined in steps 3(c) and 3(d) are complex and unusable for the next iteration.

Various methods are described in the literature that compensate for this obvious weakness. For a systematic investigation of the methods presented below, we refer the reader to [36], which is based on empirical results.

1. Absorption (A): Use the positive part of the predecessor of the variance iteration $v((j-1)\Delta)^+$ to approximate the variance

$$v(j\Delta) = v((j-1)\Delta)^+ + \kappa(\theta - v((j-1)\Delta)^+)\Delta + \sigma\sqrt{v((j-1)\Delta)^+\Delta}Z_3$$

and to determine $\mathcal{X}(j\Delta)$ in the simulation step.

2. Reflection (R): Use the absolute amount of the predecessor of the variance iteration, that is,

$$v(j\Delta) = |v((j-1)\Delta)| + \kappa(\theta - |v((j-1)\Delta)|)\Delta + \sigma\sqrt{|v((j-1)\Delta)|\Delta}Z_3$$

to determine the variance value. Use the absolute amount also for $\mathcal{X}(j\Delta)$ in the simulation step.

3. Higham and Mao (HM): Use the absolute amount $|v((j-1)\Delta)|$ only in the root terms, that is, once each for calculating the succeeding value of the variance and the stock price. The other expressions of $v((j-1)\Delta)$ remain unchanged (see [30]).
4. Partial truncation (PT): Use the positive part $v((j-1)\Delta)^+$ of the preceding value of the variance approximation only in the root terms to calculate the succeeding value of the variance and the stock price. The other incidences of $v((j-1)\Delta)$ remain unchanged (see [24]).
5. Full truncation (FT): Use the positive part of the predecessor of the variance iteration in the drift and diffusion component of the variance approximation, that is,

$$v(j\Delta) = v((j-1)\Delta) + \kappa(\theta - v((j-1)\Delta)^+)\Delta + \sigma\sqrt{v((j-1)\Delta)^+\Delta}Z_3$$

and for $\mathcal{X}(j\Delta)$ in the simulation step (see [36]).

Table 3 Simulated prices and standard deviations (in parentheses) for a European call in the Heston model with $S_0 = K = 100$, $T = 1$, $\nu_0 = \theta$, $r = 0.05$, and $\rho = -0.9$. Moreover, $n = 100000$. The exact value corresponds to the analytical price

κ	θ	σ	exact	A	R	HM	PT	FT
2.0	0.04	0.2	9.349	9.328 (0.037)	9.328 (0.037)	9.328 (0.037)	9.328 (0.037)	9.328 (0.037)
		0.5	8.881	8.890 (0.029)	8.919 (0.029)	8.871 (0.029)	8.890 (0.029)	8.859 (0.029)
	0.01	0.2	5.594	5.584 (0.017)	5.585 (0.017)	5.582 (0.017)	5.584 (0.017)	5.582 (0.017)
		0.5	5.156	5.512 (0.014)	5.862 (0.016)	5.593 (0.019)	5.516 (0.014)	5.149 (0.012)
0.5	0.04	0.2	9.278	9.255 (0.034)	9.255 (0.034)	9.255 (0.034)	9.255 (0.034)	9.255 (0.034)
		0.5	8.317	8.467 (0.024)	8.633 (0.025)	8.534 (0.027)	8.468 (0.024)	8.307 (0.023)
	0.01	0.2	5.507	5.515 (0.014)	5.533 (0.015)	5.515 (0.014)	5.515 (0.014)	5.496 (0.014)
		0.5	4.723	5.352 (0.012)	6.041 (0.016)	5.852 (0.027)	5.355 (0.012)	4.743 (0.009)

The example in Table 3 confirms the result in [36]; namely, of all the methods described above, full truncation functions best. Here, we consider a European call with a residual term of one year, and use $n = 100000$ paths for the Monte Carlo simulation. The remaining parameters are chosen so that, for falling κ and θ and rising σ , the discretized variance process becomes more frequently negative and the various truncation methods must be applied. In the table, we present the analytical value, the simulated option price, and, in parentheses, the standard deviation of the option price estimator.

One notices here—especially in cases where the variance process must be modified frequently—considerable price differences for similarly small and therefore unremarkable standard deviations. Thus, the danger for practical application is that incorrect option prices having small standard deviations might mistakenly be considered good. All told, we can propose the following simple procedure:


1. Repeat for $i = 1, \dots, N$:
 - (a) Simulate one path each of the Heston price process $S(t)$ and the variance process $\nu(t)$, $t \in [0, T]$ as described above, using the Euler–Maruyama scheme and the FT variant for the variance process.
 - (b) Calculate the corresponding option payoff $Y^{(i)}$.
2. Estimate the option price X_Y as

$$X_Y := e^{-rT} \frac{1}{N} \sum_{i=1}^N Y^{(i)}.$$

6.2.2 Tree Approaches

The Monte Carlo simulation technique, which simulates stock price paths successively, is especially well suited for pricing path-dependent derivatives.

Fig. 10 Example of a two-step binomial approximation




When pricing products that allow for multiple exercise times or even for a permanent possibility to exercise the option—so-called Bermuda or American options—tree methods offer simpler and more efficient approaches than Monte Carlo methods. Here, for each time increment, one assumes several possible developments—the next points or next nodes—and assigns them transition probabilities. The next nodes therefore represent possible future stock prices, each of which has a different probability. To determine the option prices, the nodes are then processed from the leaves toward the root using backward induction. Analogously to the algorithm of the Euler–Maruyama method for path generation, $\Delta = T/n$.

To approximate efficiently, it is crucial to be able to calculate both the nodes of the tree and the transition probabilities before the actual backward induction. Moreover, if the probabilities are chosen so that the first two moments of the price increments of the continuous and the approximated models coincide, then, according to Donsker's theorem, the tree approximation converges to the continuous process. For a detailed examination of both the standard approximation methods in the Black–Scholes model and the theory of convergence, we refer the reader to [4]. Figure 10 shows a two-period binomial tree.

We now present an algorithm developed at the ITWM by Ruckdeschel, Sayer, and Szimayer (see [9]) that achieves an efficient tree approximation in the Heston model. The method's fundamental idea is to model the variance and stock price processes as separate trees and to incorporate the correlation of the Brownian motions via a modification of the resulting transition probabilities.

However, because the variance process is mean stationary and its diffusion component depends on the current value $v(t)$, a naive approximation of the process leads to difficulties during implementation. The tendency to revert to the mean causes the process drift to become larger as the process moves further away from θ . For large trends, however, one sees negative and thus non-admissible transition probabilities. On the other hand, the dependency of the diffusion component on the current state leads to jump heights that depend on the starting level.

Fig. 11 Binomial approximation of the variance process



For such a tree approximation, the number of nodes increases exponentially, the computational effort increases, and the tree becomes inefficient, that is, useless for practical application. Here, the Itô transformation

$$R(t) = \frac{2\sqrt{v(t)}}{\sigma},$$

offers a remedy, since the variance of the resulting process

$$dR(t) = \left(\left(\frac{2\kappa\theta}{\sigma^2} - \frac{1}{2} \right) \frac{1}{R(t)} - \frac{\kappa}{2} R(t) \right) dt + dW_2(t), \quad R(0) = \frac{2\sqrt{v_0}}{\sigma}$$

is constant and a binomial approximation re-combines, since all approximation nodes exhibit the distance $\sqrt{\Delta}$. Inversion of the transformation allows one to then determine the variance values for the detected nodes. If, for each of these nodes, one now chooses successors that surround the drift, one can ensure that the transition probabilities are positive and add up to one, and that the approximation converges to the continuous model.

Figure 11 shows a variance approximation. Note, first, that the state-dependent diffusion causes the node intervals to increase as one moves upward and, second, that one sees irregular jumps—that is, jumps with multiple jump heights—for small variances, due to the tendency to revert to the mean.

To approximate the stock price, [9] uses a trinomial tree. Although this increases the computational effort, it also improves the accuracy of the approximation. Analogously to the variance approximation, the diffusion component of the logarithmized stock price process is not constant, but depends instead on the current variance value, that is, on the current node of the variance approximation. Therefore, a naive approximation also leads here to a non-efficient (from a numerical perspective) tree.

One possible way around this problem is to define a constant \tilde{v} , which describes the smallest variance unit allowed for the approximation. Possible approximation nodes then exhibit the distance $\sqrt{\tilde{v}\Delta}$. If one also defines all needed stock jumps as integer multiples of this unit, then the nodes of the stock price approximation lie on a uniform grid and

the approximation re-combines. In order to ensure convergence, one determines the transition probabilities in the model such that the first two moments in the continuous and approximated models coincide.

In summary, with this approach, one has determined the tree approximation for the variance and stock price processes, since the node set, each successor node, and the transition probabilities are known.

The next step is to combine both separate approximations into one tree model. Here, one must determine the successor nodes and transition probabilities for each possible combination of the two node sets.

The successor nodes for a node combination of the stock and variance approximations are given by the six combinations of each successor node of the separate approximations. For uncorrelated Brownian motions, each transition probability is calculated as the product of the separate probabilities. For a non-zero correlation, the authors of [9] introduce an adjustment of the product probabilities that retains the marginal moments already determined in the course of preparing the separate trees. Because the adjustment of the probabilities can be determined before the actual backward induction and the tree approximation is re-combining, the resulting approximation method is fast and accurate, even for high correlation values.

[13] presents an application of the algorithm described here. In this application, the author prices employee stock options having permanent exercise rights and specific execution hurdles.

6.3 The Complex Logarithm—An Important Detail for Implementation

In financial mathematics, the use of characteristic functions for product pricing is based, in particular, on the very generally applicable price formula from [22], which is, in turn, based on a fast Fourier transformation. This representation also forms the theoretical basis of Theorem 4 for analytical solutions in the Heston model. Implementing and numerically evaluating this semi-closed formula requires the use of complex values, which for our purposes, are incorrectly treated under some circumstances.

In order to permit a detailed investigation of the problem in the following discussion, we waive the case distinction from Theorem 4 by defining

$$\varphi(u) = \varphi_2(S(t), v(t), t, u)$$

and using the relationship

$$\varphi_1(S(t), v(t), t, u) = \frac{e^{-r(T-t)}}{S(t)} \varphi(u - i).$$

The characteristic function then becomes

Fig. 12 Real part of the complex logarithm for a complex number z





Fig. 13 Imaginary part of the complex logarithm for a complex number z



$$\begin{aligned} \varphi(u) &= \exp(iu(\ln(S(t)) + r\tau)) \\ &\times \exp\left(\frac{\kappa\theta}{\sigma^2}\left((\kappa - \rho\sigma ui + d)\tau - 2\ln\left(\frac{1 - ge^{d\tau}}{1 - g}\right)\right)\right) \\ &\times \exp\left(\frac{v(t)}{\sigma^2}(\kappa - \rho\sigma ui + d)\frac{1 - e^{d\tau}}{1 - ge^{d\tau}}\right) \end{aligned}$$

with

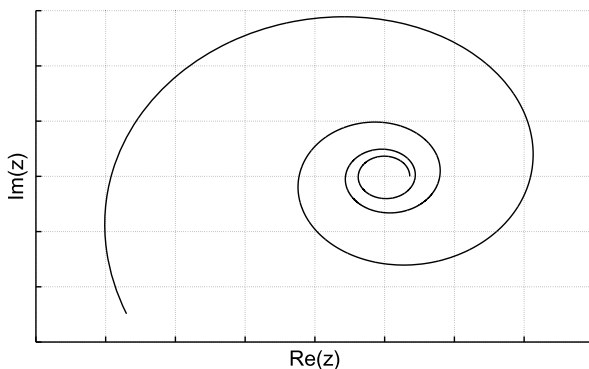
$$\tau = T - t, \quad g = \frac{\kappa - \rho\sigma ui + d}{\kappa - \rho\sigma ui - d}$$

and

$$d = \sqrt{(\rho\sigma ui - \kappa)^2 + \sigma^2(ui + u^2)}. \tag{41}$$

A significant problem with the implementation is the complex logarithm, which, in contrast to a real logarithm, is not unique. The standard software systems used for pricing financial products typically implement the principal value of the complex logarithm. Figures 12 and 13 show the real and imaginary parts of the complex logarithm for different branches.

Fig. 14 Trajectory of $(1 - ge^{d\tau})/(1 - g)$ in the complex plane



Due to the non-continuity described earlier, the integration of the characteristic function—which must be performed to determine the price in Theorem 4—is not stable starting at a certain residual time-to-maturity τ . Frequently, the problem of the ambiguity leads to large price differences that are hard to locate as numerical difficulties. Numerical problems automatically arise for sufficiently large residual time-to-maturity if the Heston parameters are chosen such that $\kappa\theta \neq m\sigma^2$ for an integer m (see [14]). This is because the trajectory of $(1 - ge^{d\tau})/(1 - g)$ describes a spiral around the origin with an exponentially increasing radius (see Fig. 14).

If the residual time-to-maturity is large enough, the trajectory inevitably crosses the negative real axis, thus producing a discontinuity. One remedy is to add 2π to the imaginary part of the result for each crossing of the negative real axis. A more elegant variant is to modify the characteristic function. To do so, one takes

$$\begin{aligned} \tilde{\varphi}(u) &= \exp(iu(\ln(S(t)) + r\tau)) \\ &\times \exp\left(\frac{\kappa\theta}{\sigma^2}\left((\kappa - \rho\sigma ui - d)\tau - 2\ln\left(\frac{1 - \tilde{g}e^{-d\tau}}{1 - \tilde{g}}\right)\right)\right) \\ &\times \exp\left(\frac{v(t)}{\sigma^2}(\kappa - \rho\sigma ui - d)\frac{1 - e^{-d\tau}}{1 - \tilde{g}e^{-d\tau}}\right) \end{aligned}$$

with

$$\tilde{g} = \frac{\kappa - \rho\sigma ui - d}{\kappa - \rho\sigma ui + d} = \frac{1}{g}$$

as the modified characteristic function. The only difference between $\tilde{\varphi}$ and φ is the negative sign of d , that is, the choice of the negative root in Eq. (41). Since

$$\begin{aligned} d\tau - 2\ln\left(\frac{1 - ge^{d\tau}}{1 - g}\right) &= d\tau - 2\ln(e^{d\tau}) - 2\ln\left(\frac{1 - e^{-d\tau}/g}{1 - 1/g}\right) \\ &= -d\tau - 2\ln\left(\frac{1 - \tilde{g}e^{-d\tau}}{1 - \tilde{g}}\right) \end{aligned}$$

Table 4 Initial values of the calibration and calibrated parameters

	ν_0	κ	θ	σ	ρ
Initial value	0.12	3.00	0.09	0.10	-0.95
Calibration	0.28	1.01	0.21	1.50	-0.79

and

$$\frac{d(1 - e^{d\tau})}{1 - ge^{d\tau}} = \frac{d(1 - e^{-d\tau})}{g - e^{-d\tau}} = \frac{-d(1 - e^{-d\tau})}{1 - \tilde{g}e^{-d\tau}}$$

are valid, $\tilde{\varphi}$ is equivalent to φ . The trajectory of $(1 - \tilde{g}e^{-d\tau})/(1 - \tilde{g})$, however, does not cross the real negative axis and the modification $\tilde{\varphi}$ is thus more stable numerically. To implement the analytical solution, it is therefore advisable to use the characteristic function $\tilde{\varphi}$.

6.4 Empirical Quality of the Heston Model

In this section, we want to illustrate the empirical quality of the Heston model, that is, its ability to replicate reality, by calibrating a real volatility surface. The stock of Allianz SE will serve as our example.


The corresponding volatility surface from 14 December 2011, obtained from the implied volatilities of European calls, is shown in Fig. 5 in Sect. 4.3.1. The shape of the surface is characteristic for volatility surfaces in general. Thus, for a fixed maturity, an option's implicit volatility is typically lower, the closer the strike lies to the current stock price. One also observes that, for a fixed exercise price, the implied volatility declines as the term of the option increases.

The freely chosen initial values for the calibration of the Heston parameters and the calibrated results obtained by applying a deterministic minimization algorithm are listed in Table 4. Figure 15 shows the calibrated surface that results when the implicit volatilities for given maturities and execution prices are calculated and presented with the help of the calibrated Heston parameters. Typically, the calibrated surface is considerably smoother than the original, but the characteristics of the real volatility surface are retained.

7 Mathematical Modeling and Algorithmic Implementation in the Financial Market—A Few Closing Remarks

This example of option pricing in connection with the Heston model is but one of many similar research and implementation projects that have been successfully dealt with by the Financial Mathematics Department of the Fraunhofer ITWM in cooperation with partners from the financial and insurance industries. Some examples of other projects involving innovative in-house developments and algorithmic implementations are:

Fig. 15 Calibrated implied volatility surface



- development of a new stock price model based on the explicit modeling of future dividend payments, in cooperation with the University of Cambridge (see [8]);
- development of a dynamic mortality model for evaluating longevity bonds, together with the Hypovereinsbank (see [7]);
- algorithmic implementation of robust statistics in the field of operational risk (see [31]), honored with a “best paper award”;
- development of a completely new approach for efficient, multi-dimensional binomial trees (see [6]).

In addition, many of the algorithms used in the daily work of the ITWM are described extensively in [5].

There are several components common to all these projects and developments that are typical for implementations in the financial and insurance markets:

1. In general, the methods used are based on continuous-time stochastic processes and require thorough training in the fields of Itô calculus, martingale theory, and stochastic processes.
2. The client’s wish for the best possible explanation of observed market prices leads to a wish for the generalization of existing models. Here, one must always make sure that the introduction of further parameters (e.g., by replacing a constant with a deterministic function) does not lead to numerical or statistical instability.
3. The use of a variety of numerical methods (e.g., Monte Carlo simulation, tree methods, Fourier transformation) is necessary in order to be able to calculate the prices of the diverse (exotic) options. Here, the character of the option determines the choice of the numerical algorithm. There is no universal, standard algorithm that performs well for all option types.
4. Calibration of the parameters plays a very significant role. While it’s true that no spectacular theoretical results can be achieved in this domain, reliably calibrated parameters form the basis of all mathematical modeling and calculation that succeeds in the market.

5. Theoretical understanding of the models is indispensable if one is to calculate those values that are actually desired. The lack of understanding in shifting between the risk-neutral and the physical model worlds, in particular, is a frequent source of error.

Finally, it is important to emphasize the responsibility of the financial mathematician to help ensure a correct—and above all, wise—application of his models. Particularly with a view toward the financial crisis of these recent years, the financial mathematician must

- warn against mistaking the model for the reality,
- point out the inability of most models to predict, and
- avoid bringing excessive complexity into derivative products.

It was precisely the successful mathematical treatment of ever newer and more complex problems in the financial market that encouraged product designers to offer ever more complexly structured products—products whose effects were, in large measure, incomprehensible to customers but were bought anyway, despite this lack of understanding. Here too, the financial mathematician has a responsibility to warn against such dangerous developments.

References

Publications of the Authors

1. Korn, R.: *Optimal Portfolios*. World Scientific, Singapore (1997)
2. Korn, R.: *Elementare Finanzmathematik*. *Berichte Fraunhofer ITWM* **39**, 1–89 (2002)
3. Korn, R.: *Optimal portfolios—New variations of an old theme*. *Comput. Manag. Sci.* **5**, 289–304 (2008)
4. Korn, R., Korn, E.: *Optionsbewertung und Portfolio-Optimierung*. Vieweg, Wiesbaden (2001)
5. Korn, R., Korn, E., Kroisandt, G.: *Monte Carlo Methods and Models in Finance and Insurance*. Chapman & Hall/CRC Financial Mathematics Series. CRC Press, London (2010)
6. Korn, R., Müller, S.: *The decoupling approach to binomial pricing of multi-asset options*. *J. Comput. Finance* **12**, 1–30 (2009)
7. Korn, R., Natcheva, K., Zipperer, J.: *Langlebigkeitsbonds: Bewertung, Modellierung und Anwendung auf deutsche Daten*. *Blätter der DGVFM* **27**(3), 397–418 (2006)
8. Korn, R., Rogers, L.: *Stocks paying discrete dividends: modelling and option pricing*. *J. Deriv.* **13**(2), 44–49 (2005)
9. Ruckdeschel, P., Sayer, T., Szimayer, A.: *Pricing American options in the Heston model: a close look at incorporating correlation*. *J. Deriv.* **20**(3), 9–29 (2013)

Dissertations in the Area at Fraunhofer ITWM

10. Horsky, R.: *Barrier Option Pricing and CPPI-Optimization*. Ph.D. thesis, TU, Kaiserslautern (2012)

11. Krekel, M.: Some New Aspects of Optimal Portfolios and Option Pricing. Ph.D. thesis, TU, Kaiserslautern (2003)
12. Natcheva, K.: On Numerical Pricing Methods of Innovative Financial Products. Ph.D. thesis, TU, Kaiserslautern (2006)
13. Sayer, T.: Valuation of American-style derivatives within the stochastic volatility model of Heston. Ph.D. thesis, TU, Kaiserslautern (2012)

Further Literature

14. Albrecher, H., Mayer, P., Schoutens, W., Tistaert, J.: The little Heston trap. *Wilmott Magazine*, 83–92 (2007)
15. Artzner, P., Delbean, F., Eber, J.M., Heath, D.: Coherent measures of risk. *Math. Finance* **9**(3), 203–228 (1999)
16. Bachelier, L.F.: Théorie de la spéculation. *Ann. Sci. Éc. Norm. Super.* **17**, 21–86 (1900)
17. Barndorff-Nielsen, O., Shephard, N.: Non-Gaussian Ornstein–Uhlenbeck based models and some of their uses in financial economics. *J. R. Stat. Soc. B* **63**, 167–241 (2001)
18. Bingham, N., Kiesel, R.: Risk-Neutral Valuation: Pricing and Hedging of Financial Derivatives. Springer, Berlin (1998)
19. Björk, T.: Arbitrage Theory in Continuous Time, 2nd edn. Oxford University Press, Oxford (2004)
20. Black, F., Scholes, M.: The pricing of options and corporate liabilities. *J. Polit. Econ.* **81**, 637–654 (1973)
21. Broadie, M., Kaya, Ö.: Exact simulation of stochastic volatility and other affine jump diffusion processes. *Oper. Res.* **54**(2), 217–231 (2006)
22. Carr, P., Madan, D.: Option valuation using the fast Fourier transform. *J. Comput. Finance* **2**, 61–73 (1999)
23. Cont, R., Tankov, P.: Financial Modelling with Jump Processes. Financial Mathematics Series. Chapman & Hall, Boca Raton (2003)
24. Deelstra, G., Delbaen, F.: Convergence of discretized stochastic (interest rate) processes with stochastic drift term. *Appl. Stoch. Models Data Anal.* **14**(1), 77–84 (1998)
25. Delbaen, F., Schachermayer, W.: The Mathematics of Arbitrage. Springer, Berlin (2006)
26. Dimitroff, G., Lorenz, S., Szimayer, A.: A parsimonious multi-asset Heston model: calibration and derivative pricing. *Int. J. Theor. Appl. Finance* **14**(8), 1299–1333 (2011)
27. Dupire, B.: Pricing and hedging with smiles. In: Dempster, M.A., Pliska, S.R. (eds.) *Mathematics of Derivative Securities*, pp. 103–111. Cambridge University Press, Cambridge (1997)
28. Eberlein, E., Keller, U.: Hyperbolic distributions in finance. *Bernoulli* **1**, 281–299 (1995)
29. Heston, S.L.: A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. Financ. Stud.* **6**(2), 327–343 (1993)
30. Higham, D.J., Mao, X.: Convergence of Monte Carlo simulations involving the mean-reverting square root process. *J. Comput. Finance* **8**(3), 35–61 (2005)
31. Horbenko, N., Ruckdeschel, P., Bae, T.: Robust estimation of operational risk. *J. Oper. Risk* **6**, 3–30 (2011)
32. Jäckel, P.: Monte Carlo Methods in Finance. Wiley, West Sussex (2002)
33. Karatzas, I., Shreve, S.E.: Brownian Motion and Stochastic Calculus, 2nd edn. Springer, Berlin (1991)
34. Kloeden, P.E., Platen, E.: Numerical Solution of Stochastic Differential Equations. Springer, Berlin (1999)

35. Kruse, S., Nögel, U.: On the pricing of forward starting options in Heston's model on stochastic volatility. *Finance Stoch.* **9**, 233–250 (2005)
36. Lord, R., Koekkoek, R., van Dijk, D.: A comparison of biased simulation schemes for stochastic volatility models. *Quant. Finance* **2**, 177–194 (2010)
37. Madan, D.B., Seneta, E.: The variance gamma model for share market returns. *J. Bus.* **63**, 511–524 (1990)
38. Mikhailov, S., Nögel, U.: Heston's Stochastic Volatility Model Implementation, Calibration and Some Extensions. *Wilmott Magazine* (2003)
39. Mishra, S.: Optimal Solution of the nearest Correlation Matrix Problem by Minimization of the Maximum Norm (2004). <http://mpa.ub.uni-muenchen.de/1783/>
40. Schoutens, W.: *Lévy Processes in Finance: Pricing Financial Derivatives*. Wiley, New York (2003)