

S Plus
For Financial Engineers ¹
-
Part IV
-
Pricing and Hedging
of Options

Diethelm Würtz

Institut für Theoretische Physik
ETH Zürich

June 26, 2002

¹This script collects material used in my lecture *Econophysics* held in SS 2002 at the "Institute of Theoretical Physics" of ETH Zürich. These notes are thought for internal use only, please do not distribute! The software used in the script comes without any warranty!

Overview

Chapter 1 - Markets, Basic Statistics, Date and Time

- 1.1 Introduction
- 1.2 Economic and Financial Markets
- 1.3 Brief Repetition from Probability Theory
- 1.4 Distribution Functions in Finance
- 1.5 Searching for Correlations and Dependencies
- 1.6 Hypothesis Testing
- 1.7 Calculating and Managing Calendar Dates

Chapter 2 - The Dynamical Process Behind Financial Markets

- 2.1 Introduction
- 2.2 ARIMA Modelling: Basic Concepts of Linear Processes
- 2.3 GARCH Modelling: Mastering Heteroskedastic Processes
- 2.4 Regression Modelling from the Time Series Point of View
- 2.5 Neural Networks: Feedforward Connectionist Networks
- 2.6 Design and Implementation of an Intra-Daily Trading System

Chapter 3 - Beyond the Sample: Dealing With Extreme Values

- 3.1 Introduction
- 3.2 Fluctuations of Maxima
- 3.3 Extremes of Point Processes
- 3.4 The Extremal Index

Chapter 4 - Pricing and Hedging of Options

- 4.1 Introduction
- 4.2 The Basics of Option Pricing
- 4.3 Pricing Formulas for Exotic Options
- 4.4 Heston Nandi Option Pricing
- 4.5 MC Simulation of Path Dependent Options

Appendix

An Introduction to Splus

Contents

4 Pricing and Hedging of Options	6
4.1 Introduction	6
4.2 The Basics of Option Pricing	7
4.2.1 Black-Scholes Option Pricing	12
4.2.2 Options Sensitivities	21
4.2.3 Analytical Pricing Formulas for American Options	28
4.2.4 Binomial Option Pricing	30
4.3 Pricing Formulas for Exotic Options	34
4.3.1 Options with Contract Variations	35
4.3.2 Simple Path Dependent Options	36
4.3.3 Limit Dependent Options	39
4.3.4 Multiple Factors Options	40
4.3.5 Pricing Formulas for Other Exotic Options	41
4.4 Heston Nandi Option Pricing	42
4.4.1 The Heston Nandi GARCH Option Pricing Model	43
4.4.2 Numerical Analysis of the Heston Nandi Model	49
4.5 Monte Carlo Simulations of Options	64
4.5.1 The Monte Carlo Approach	64
4.5.2 Monte Carlo Estimators of the Greeks	66
4.5.3 Variance-Reduction Techniques	67
4.5.4 Monte Carlo Importance Sampling	76
4.6 The <code>fOptions</code> Library	84
4.6.1 Summary of Splus Funcions	84
4.6.2 List of Datasets	85
4.6.3 List of Splus Examples	86
4.6.4 Software Packages	86

Chapter 4

Pricing and Hedging of Options

4.1 Introduction

In this Chapter we present methods for pricing and hedging options.

Section 2 introduces the basics of option prices. We derive the Black-Scholes option pricing formula and discuss options sensitivities. We also present some formulas for pricing analytically American Options and briefly discuss binomial option pricing.

Section 3 is devoted to pricing exotic options. First we consider options with contract variations. Then we present results for path dependent options and for multiple factors options. Finally we summarize some pricing formulas for other exotic options.

Section 4 deals with the Heston-Nandi Option Pricing approach. We discuss the analytical results and show how to solve numerically the results.

Section 5 is concerned with Monte Carlo simulations in the field of options pricing. We introduce to the Monte Carlo approach, present Monte Carlo estimators of the Greeks, discuss the case of path-dependent options, and show how variance reduction techniques work. Additionally we discuss the usage of low discrepancy sequences for random numbers and close the section with an application to option data from the Sydney Futures Exchange.

4.2 The Basics of Option Pricing

Introduction

A derivative security is a security whose value depends on the values of other more basic underlying variables. In recent years derivative securities have become increasingly important in the field of finance. Futures and options are now actively traded on many different exchanges, and outside of exchanges forward contracts swaps, and many different types of options are traded by financial institutions and their corporate clients in what are termed over-the-counter markets.

This section will give an intuitive introduction to options and wants to introduce into the numerical approaches of pricing derivatives.

What are Options?

Options are derivatives which imply different types of rights and obligations, expressed as Calls and Puts. These Calls and Puts are traded according different rules depending on their special contract specifications. To become more specific we will give short definitions of Calls, Puts, Premium, European/Asian Options, Traded/OTC Options, Warrants, Plain Vanilla/Exotic Options, and types of Traders.

Calls: Calls give the buyer the right, but not the obligation, to *buy* a given quantity of the underlying asset, at a given price, known as the *exercise price* or *strike price*, on or before a given future date, the *maturity date* or *expiry date*.

Puts: Puts give the buyer the right, but not the obligation, to *sell* a given quantity of the underlying asset, at a given price on or before a given future date.

Premium: New options are created by one party selling them and thereby undertaking the obligations embedded in the options contract. Unlike the buyer, the seller known as the *writer* has no choice regarding the fulfillment of the obligations under the options contract. If the buyer wants to exercise his right, the writer must comply. For this asymmetry of privilege the buyer must pay the writer the option price what is known as the *premium*.

European / American Options: Options are also classified as *European* or *American*. With European options the right to buy or sell the underlying asset can only be exercised on the expiry date, but with American options that right can be exercised at any time.

Traded / OTC Options: Traded options are those options traded on recognized *exchanges* and where there is an active *secondary market*. Options are also sold by private negotiation; they are said to be traded *over-the-counter* in the so called *OTC markets*. Traded options are standardized contracts, traded according to the rules of the particular options exchange. OTC options, on the other hand, are individually tailored to the needs of the customer.

Warrants: Longer traded options are called *warrants*, these too can be American or European, and are generally traded over-the-counter. However, some are listed at stock exchanges but generally not on the recognized option exchanges.

Plain Vanilla / Exotic Options: The options so far mentioned are also called *standard* or *plain vanilla* options. Options with special properties attached, are named exotic options. One kind of this class of options are path dependent options; barrier options, look-back options, binary options, and many other types of options. They can be of European or American style.

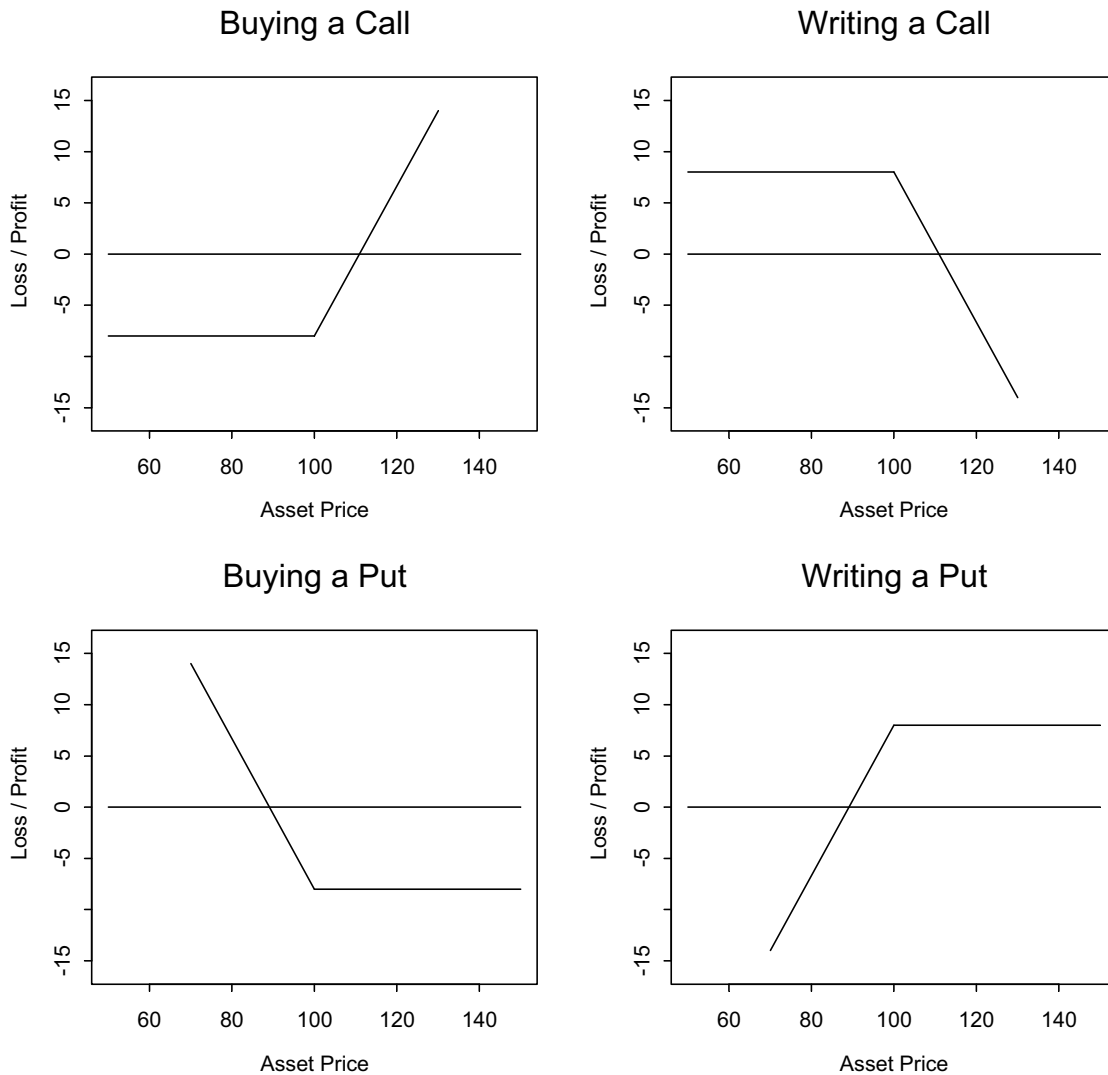
Types of Traders: Traders of derivative securities can be categorized as hedgers, speculators, or arbitrageurs. *Hedgers* are interested in reducing a risk that they already face. Whereas hedgers want to eliminate an exposure to movements in the price of an asset, *speculators* wish to take a position in the market. Either they are betting that a price will go up or they are betting that it will go down. *Arbitrageurs* are a third important group of participants in derivative security markets. Arbitrage involves locking in a risk-less profit by simultaneously entering into transactions in two or more markets.

Buying and Selling Options

To get more familiar of the principles of buying and selling options we will follow the four basic operations: buying a call, writing a call, buying a put, and writing a put.

Buying a Call: The call buyer pays the option premium of 8 in return for the right to buy the underlying asset at the exercise price, X , of 100. If at the expiry date of the option the underlying asset price, S , is above the exercise price, say 120, the buyer will exercise the option, pay the exercise price, and receive the asset. This may then be sold in the market at 120 giving an instant profit of $12 = (120 - 100) - 8$. Alternatively, the option may be sold prior to expiry to realize a similar profit, because at expiry, its value must be equal to the difference between the exercise price and the market price of the underlying asset, otherwise arbitrage profits would be possible. If, on the other hand, the asset price at expiry is at or below the exercise price, the option will be abandoned by the buyer and he/she will suffer a loss equal to the option premium.

Writing a Call: The option writer is paid the option premium of 8 as compensation of bearing the risk of having to deliver the underlying asset in return for being paid the exercise price. If at the expiry, the asset price is, say 120, i.e. above the exercise price of 100, the writer will incur a loss because he/she will have to buy the asset at the market price in order to deliver to the options buyer in exchange for the lower exercise price. In this example the loss will be $(120 - 100) - 8 = 12$. If however, the asset price is below the exercise price at expiry, say 90, the call option will not be exercised and the writer will make a profit equal to the option premium.



■ Figure 4.1.1: The two upper graphs show the payoff functions for buying (left) and writing (right) a call. The two lower graphs show the payoff functions for buying (left) and writing (right) a put. For an explanation we refer to the text.

Buying a Put: The buyer pays the option premium of 8 for the right to sell (put) the underlying asset at the exercise price of 100. If at expiry, the asset price is below the exercise price, say 90, the put will have a value equal to the difference between the two prices, i.e. $100 - 90 = 10$, and the option will show a profit of $10 - 8 = 2$. If however, the asset price is at or above the exercise price, the put option will be abandoned and the buyer will incur a loss equal to the option premium.

Writing a Put: The put writer is paid a premium of 8 for bearing the risk of having to take the underlying asset at the exercise price. If the market price of the asset is below the exercise price at expiry, the writer will incur a loss because he/she will have to pay the exercise price, 100, but will only be able to resell the asset at the lower market price, i.e. 90. If

however, the asset price is above the exercise price at expiry, the buyer will abandon the put option and the writer will make a profit equal to option premium.

Factors Influencing Option Prices

Clearly at expiry, only two factors influence the value of an option, the exercise price and the asset price. If the option has time remaining to expiry a number of factors come into play:

The Asset Price and Exercise Price: The higher the asset price S , relative to the exercise price X , the more valuable a call will be, and the less valuable will be a put. In the case of a call, where the difference $X - S > 0$ the option is said to have *intrinsic value* and to be *in-the-money*. For a put to have intrinsic value, i.e. to be *in-the-money*, the asset price must be below the exercise price. If, $X = S$, the options are said to be *at-the-money*. In the case of a call, if $S < X$, the option is said to be *out-of-the-money*. It will have no intrinsic value and what value it does have - known as *time value* or *extrinsic value* - is dependent upon the interaction of the other factors influencing the option value. A put is out-of-the money if $S > X$, in which case any value it does have will be extrinsic value. The term time value or extrinsic value is used to describe that part of an option premium that is not represented by intrinsic value. For example, if an option is priced at 20 with the exercise price at 240 and the security price of 255, the intrinsic value would be $15=255-240$, and the time value would be $5=20-15$.

The Time to Maturity: The longer the time to maturity $T - t$, the greater is the probability that the asset price S will be substantially different from the exercise price X and, as this probability has some utility, the higher will be the value of both puts and calls.

The Rate of Interest: The interest rate r influences call options values because, by buying the option and not the underlying security, the buyer is making a highly geared investment thereby releasing capital to be invested at the risk free rate. There is an opportunity saving in buying a call option and not the asset, this saving being higher, the higher the rate of interest. Therefore, the higher the rate of interest, the higher the value of the call option. This may also be expressed by saying that the higher the rate of interest the lower will be the present value of the exercise price and, therefore, the higher will be the value of the call. However, in the case of a put option, the lower the present value of the exercise price, the lower will be the value of the option. Thus higher interest rates will result in lower put option prices.

The Volatility of the Underlying Security: The more volatile the underlying asset price, the more valuable will be the option. This is because the greater the volatility, the greater the probability of the asset price changing substantially to be above (for a call) or below (for a put) the exercise price at expiry. It is true that the asset price may also fall (rise), but, as the option does not have to be exercised, the adverse movement is avoided. Thus the greater the volatility, the greater the potential gain for the option holder and, both calls and puts will be more valuable.

To summarize, a call option premium will be higher, the higher the asset price relative to the exercise price, the higher the volatility, the higher the rate of interest and the longer the time to maturity. Put option premiums will be higher, the lower the asset price relative to the exercise

price, the higher the volatility, the lower the rate of interest and the longer the term to maturity.

The Basic Idea Behind Option Pricing

Suppose the current price of an asset is $S = 50$, and at the end of a period of time, its price must be either $S^* = 25$ or $S^* = 100$. A call on asset is available with an exercise price of $X = 50$, expiring at the end of the period. It is also possible to borrow and lend at a $r = 25\%$ rate of interest.

The one piece of information left unfurnished is the current value of the call, c . However, if profitable risk-less arbitrage is not possible, we can deduce from the given information *alone* what the value of the call *must* be!

Example: A Leveraged Hedge

Consider forming the following leveraged hedge:

1. Write (sell) three calls at c each
2. Buy two shares at 50 each
3. Borrow 40 at 25% to be paid back at the end of the period

The return from this hedge for each possible level of the asset price at the expiration date is:

	Current Date $X = 50; r = 25\%$		Expiration Date $S^* = 25$	Expiration Date $S^* = 100$
Write 3 calls	3 c	Payoff	-	-150
Buy 2 shares	-100	Deliver	50	200
Borrow	40	Pay Back	-50	-50
Total			0	0

Regardless of the outcome, the hedge exactly breaks even on the expiration date. Therefore, to prevent profitable risk-less arbitrage, the current cash flow from establishing the position must be zero; that is,

$$3c - 100 + 40 = 0.$$

The current value of the call must be $c = 20$. If the call were not priced at 20, a sure profit would be possible. In particular if $c = 25$, the hedge would yield a current amount of 15 and would experience no further gain or loss in the future. On the other hand, if $c = 15$, then the same thing could be accomplished by buying three calls, selling short two shares, and lending 40.

The table above can be interpreted, as demonstrating an *appropriately levered position in asset will replicate the future returns of a call*. That is, if we buy assets and borrow against them in the right proportion, we can in effect, duplicate a pure position in calls. In view of this, it should seem less surprising that all we needed to determine the *exact* value of the call was its *exercise price, underlying asset price, range of movement in the underlying asset price, and the rate of interest*. What may seem more incredible is what we do not need to know: Among other things, *we do not need to know the probability that the asset price will rise or fall*.

Although this example was quite simple, it showed several essential features of option pricing. And we will see later that it was not so unrealistic as it seemed.

4.2.1 Black-Scholes Option Pricing

Wiener Processes - Brownian Motion

Standard models of asset or stock price behavior are expressed in terms of *Wiener processes*, also referred to under the name *Brownian motion*.

The behavior of a variable z , which follows a Wiener process, can be understood by considering the changes in its value in small intervals of time. Consider a small interval of time Δt and define Δz as the change in z during Δt . There are two basic properties Δz must have for z to be following a Wiener process:

- PROPERTY 1: Δz is related to z by the equation $\Delta z = \varepsilon \cdot (\Delta t)^{1/2}$, where ε is a random drawing from a standardized normal distribution $N(0, 1)$.
- PROPERTY 2: The values of Δz for any two different short intervals of time Δt are independent.

It follows from PROPERTY 1 that Δz itself has a normal distribution with mean $\Delta z = 0$, and variance of $\Delta z = \Delta t$. Additionally, PROPERTY 2 implies that z follows a Markov process, i.e. a stochastic process where only the present value is relevant for predicting the future. The past history of the variable and the way in which the present has emerged from the past are irrelevant.

Ito's Lemma

The price of stock options is a function of the underlying stock's price and time. More generally, we can say that the price of any derivative security is a function of the stochastic variables underlying the derivative security and time. Therefore we must acquire some understanding of the behavior of functions of stochastic variables. An important result in this area is known as Ito's Lemma (1951).

Suppose that the value of a variable x follows an Ito process:

$$dx = a(x, t)dt + b(x, t)dz , \quad (4.1)$$

where dz is a Wiener process and a and b are functions of x and t . The variable x has a drift rate of a and a variance rate of b^2 . Ito's shows that a function, G of x and t follow the process

$$dG = \left\{ \frac{\partial G}{\partial x} a + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2 \right\} dt + \frac{\partial G}{\partial x} b dz , \quad (4.2)$$

where the dz is the same Wiener process in both equations above. Thus G also follows an Ito process. It has a drift rate of

$$\frac{\partial G}{\partial x} a + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2 \quad (4.3)$$

and a variance rate of

$$\frac{\partial^2 G}{\partial x^2} b^2 . \quad (4.4)$$

If we approximate the asset price through the model

$$dS = uSdt + \sigma Sdz \quad (4.5)$$

with constant values u (mean) and σ (variance), it follows from Ito's Lemma, that the process followed by a function, G , of S and t is

$$dG = \left\{ \frac{\partial G}{\partial x} uS + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} \sigma^2 S^2 \right\} dt + \frac{\partial G}{\partial x} \sigma Sdz . \quad (4.6)$$

Note, that both G and S are affected by the same underlying source of uncertainty, dz . This proves to be very important in the derivation of the Black-Scholes results.

Example: Application to the Logarithm of the Asset Price

We now use Ito's Lemma to derive the process followed by $\ln S$. Define $G = \ln S$. Since $\partial G/\partial S = 1/S$, $\partial^2 G/\partial S^2 = 1/S^2$, and $\partial G/\partial t = 0$ it follows that the process followed by G is

$$dG = \left(u - \frac{1}{2} \sigma^2 \right) dt + \sigma dz . \quad (4.7)$$

Since u and σ are constant, this equation indicates that G follows a generalized Wiener process. It has constant drift rate $u - \sigma^2/2$ and constant variance rate σ^2 . This means that the change in G between the current time, t , and some future time, T , is normally distributed with mean $(u - \sigma^2/2)(T - t)$ and variance $\sigma^2(T - t)$. The value of G at time t is $\ln S$. Its value at time T is $\ln S_T$, where S_T is the asset price at time T . Its change during the time interval $T - t$ is therefore $\ln S_T - \ln S$; hence

$$\ln S_T - \ln S = \mathcal{N}[(u - \sigma^2/2)(T - t), \sigma(T - t)^{1/2}] . \quad (4.8)$$

Derivation of Ito's Lemma

A completely rigorous proof of Ito's Lemma is beyond of this course. However, Ito's Lemma can be regarded as a natural extension of other, simpler results. Consider a continuous and differentiable function G of a variable x . If Δx is a small change in x and ΔG is the resulting small change in G , it is well known that

$$\Delta G \approx \frac{dG}{dx} \Delta x . \quad (4.9)$$

In other words, ΔG is approximately equal to the rate of change of G with respect to x multiplied by Δx . The error involves terms of order Δx^2 . If a higher precision is required, a Taylor series expansion of ΔG can be used.

$$\Delta G = \frac{\partial G}{\partial x} \Delta x + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} \Delta x^2 + \frac{1}{6} \frac{\partial^3 G}{\partial x^3} \Delta x^3 + \dots . \quad (4.10)$$

For a continuous and differentiable function G of two variables x and y the result is analogous, and the Taylor series expansion is

$$\Delta G = \frac{\partial G}{\partial x} \Delta x + \frac{\partial G}{\partial y} \Delta y + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} \Delta x^2 + \frac{1}{2} \frac{\partial^2 G}{\partial x \partial y} \Delta x \Delta y + \frac{1}{2} \frac{\partial^2 G}{\partial y^2} \Delta y^2 + \dots \quad (4.11)$$

In the limit as Δx and Δy tend to zero the first order yields

$$dG = \frac{\partial G}{\partial x} dx + \frac{\partial G}{\partial y} dy . \quad (4.12)$$

A derivative security is a function of a variable that follows a stochastic process. We now extend the equation above to cover such functions. Suppose that a variable x follows the general Ito process

$$dx = a(x, t)dt + b(x, t)dz \quad (4.13)$$

and that G is some function of x and of time, t . By analogy the the Taylor series expansion we can write

$$\Delta G = \frac{\partial G}{\partial x} \Delta x + \frac{\partial G}{\partial t} \Delta t + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} \Delta x^2 + \frac{1}{2} \frac{\partial^2 G}{\partial x \partial t} \Delta x \Delta t + \frac{1}{2} \frac{\partial^2 G}{\partial t^2} \Delta t^2 + \dots \quad (4.14)$$

Discretizing the Ito process

$$\Delta x = a(x, t)\Delta t + b(x, t)\varepsilon\sqrt{\Delta t} \quad (4.15)$$

reveals an important difference between the situation in eqn. (4.14) and the situation in eqn. (4.11). When limiting arguments were used to move from eqn. (4.12) to eqn. (4.12), terms in Δx^2 were ignored because they were second order terms. From eqn. (4.15)

$$\Delta x^2 = b^2 \varepsilon^2 \Delta t + \text{higher order terms in } \Delta t \quad (4.16)$$

which shows that the term involving Δx^2 in eqn. (4.14) has a component that is of order Δt and cannot be ignored.

The variance of the standardized normal distribution is one. This means that

$$E[\varepsilon^2] - E[\varepsilon]^2 = 1 , \quad (4.17)$$

where E denotes the expected value. Since $E[\varepsilon] = 0$, it follows that $E[\varepsilon^2] = 1$. The expected value of $\varepsilon^2 \Delta t$ is therefore Δt . It can be shown that the variance of $\varepsilon^2 \Delta t$ becomes non-stochastic and equal to its expected value of Δt as Δt tends to zero. It follows that the first term to the right hand side of eqn. (4.16) becomes non-stochastic and equal to $b^2 \Delta t$ as Δt tends to zero. Taking limits as Δx and Δt tends to zero in eqn. (4.14) and using this last result, we therefore obtain

$$dG = \frac{\partial G}{\partial x} dx + \frac{\partial G}{\partial t} dt + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2 dt . \quad (4.18)$$

This is Ito's Lemma. Substituting for dx from eqn. (4.13) eqn. 4.18 becomes

$$dG = \left\{ \frac{\partial G}{\partial x} a + \frac{\partial G}{\partial t} + \frac{1}{2} \frac{\partial^2 G}{\partial x^2} b^2 \right\} dt + \frac{\partial G}{\partial x} b dz . \quad (4.19)$$

Derivation of Black-Scholes Differential Equation

We assume that the asset price S follows the process (4.7). Suppose that f is the price of a derivative security contingent on S . The variable f must be some function of S and t . From Ito's Lemma

$$df = \left\{ \frac{\partial f}{\partial S} uS + \frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2 \right\} dt + \frac{\partial f}{\partial S} \sigma S dz . \quad (4.20)$$

It follows that by choosing a portfolio of the asset and the derivative security, the Wiener process can be eliminated. The appropriate portfolio is: -1 derivative security and $+\partial f/\partial S$ shares. The holder of this portfolio is *short* one derivative security and *long* an amount of $+\partial f/\partial S$ shares. Define V as the value of the portfolio; by definition

$$V = -f + \frac{\partial f}{\partial S} S . \quad (4.21)$$

The change in the value dV of the portfolio in time dt is given by

$$dV = -df + \frac{\partial f}{\partial S} dS . \quad (4.22)$$

Substituting eqns. (4.5) and (4.20) into eqn. (4.22) yields

$$dV = \left\{ -\frac{\partial f}{\partial t} - \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2 \right\} dt . \quad (4.23)$$

Since this equation does not involve dz , the portfolio V must be riskless during time dt . The portfolio must instantaneously earn the same rate of return as other short term riskfree securities and using the proceeds to buy the portfolio; if it earned less, they could make a riskless profit by shorting the portfolio and buying riskfree securities. It follows that

$$dV = rV dt \quad (4.24)$$

where r is the riskfree interest rate. Substituting from eqns. (4.21) and (4.23) this becomes

$$\left\{ \frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2 \right\} = r \left\{ f - \frac{\partial f}{\partial S} S \right\} dt \quad (4.25)$$

so that

$$\frac{\partial f}{\partial t} + rS \frac{\partial f}{\partial S} + \frac{1}{2} \sigma^2 S^2 = rf . \quad (4.26)$$

This eqn. (4.26) is the *Black-Scholes differential equation*. It has many solutions, corresponding to all the different derivative securities that can be defined with S as the underlying variable. The particular derivative security that is obtained when the equation is solved depends on the *boundary conditions* that are used. These specify the values of the derivative security at the boundaries of possible values of S and t . In the case of European call option, the key boundary condition is

$$f = \max[0, S_T - X] . \quad (4.27)$$

In the case of a European put option, it is

$$f = \max[0, X - S_T] . \quad (4.28)$$

The Black-Scholes Pricing Formula

Black and Scholes succeeded in solving their differential equation to obtain exact formulas for the prices of European call and put options. The expected value of a European call option at maturity in a risk neutral world is

$$E[\max(0, S_T - X)] , \quad (4.29)$$

where E denotes the expected value. From the European call option price, c is the value of this discounted at the riskfree rate of interest, that is,

$$c = e^{-r(T-t)} E[\max(0, S_T - X)] . \quad (4.30)$$

Remember, that $\ln S_T$ has the probability distribution

$$\ln S_T - \ln S \sim \mathcal{N}[(u - \frac{1}{2}\sigma^2)(T - t), \sigma(T - t)^{1/2}] . \quad (4.31)$$

Evaluating the expectation value $E[\max(0, S_T - X)]$ is an application of integral calculus, yielding

$$\begin{aligned} c &= SN(d_1) - Xe^{-r(T-t)}N(d_2) , \\ d_1 &= \frac{\ln S/X + (r + \sigma^2/2)(T-t)}{\sigma(T-t)^{1/2}} , \\ d_2 &= \frac{\ln S/X + (r - \sigma^2/2)(T-t)}{\sigma(T-t)^{1/2}} = d_1 - \sigma(T-t)^{1/2} , \end{aligned} \quad (4.32)$$

and $N(x)$ is the cumulative distribution function for a standardized normal variable. The value of an European put can be calculated in a similar way; the result is

$$p = Xe^{-r(T-t)}N(-d_2) - SN(-d_1) . \quad (4.33)$$

The formula can be used as the starting point to price several kinds of options including European options on a stock with cash dividends, options on stock indexes, options on futures, and currency options. This will be shown in the following.

Use Case: European Options on a Stock with Cash Dividends

An European option on a stock that pays out one or more cash dividends during the options's life time can be priced by the Black-Scholes formula, by replacing S with S minus the present value of the dividends.

$$S - D_1e^{-rt_1} - D_2e^{-rt_2} - \dots - D_n e^{-rt_n} , \quad (t_n < T) ,$$

where D_1 is the dividend payout one, t_1 is the time to dividend the payout, and T is the time to maturity of the option.

Use Case: Options on Stock Indexes

Merton (1973) extended the Black-Scholes model to allow for a dividend yield. The model can be used to price European call and put options on a stock or stock index paying a known dividend yield equal to q ,

$$c = Se^{-qT}N(d_1) - Xe^{-rT}N(d_2) , \quad (4.34)$$

$$p = Xe^{-rT}N(-d_2) - Se^{-qT}N(-d_1) ,$$

where

$$d_1 = \frac{\ln(S/X) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}} , \quad (4.35)$$

$$d_2 = \frac{\ln(S/X) + (r - q - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T} .$$

Use Case: Options on Futures

The Black (1976) formula can be used to price European options when the underlying security is a forward or futures contract with initial price F ,

$$c = Fe^{-rT}N(d_1) - Xe^{-rT}N(d_2) , \quad (4.36)$$

$$p = Xe^{-rT}N(-d_2) - Fe^{-qT}N(-d_1) ,$$

where

$$d_1 = \frac{\ln(F/X) + (\sigma^2/2)T}{\sigma\sqrt{T}} , \quad (4.37)$$

$$d_2 = \frac{\ln(F/X) - (\sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T} .$$

Use Case: Currency Options

The Garman and Kohlhagen (1983) modified Black-Scholes model can be used to price European currency options. The model is equal to the Merton (1973) model presented earlier. The only difference is that the dividend yield is replaced by the risk-free rate of the foreign currency r_f .

$$c = Se^{-r_f T}N(d_1) - Xe^{-rT}N(d_2) , \quad (4.38)$$

$$p = Xe^{-r_f T}N(-d_2) - Se^{-qT}N(-d_1) ,$$

where

$$d_1 = \frac{\ln(S/X) + (r - r_f + \sigma^2/2)T}{\sigma\sqrt{T}} , \quad (4.39)$$

$$d_2 = \frac{\ln(S/X) + (r - r_f - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T} .$$

The Generalized Black-Scholes Formula

The general version of the Black-Scholes model incorporates the cost-of-carry term b . It can be used to price European options on stocks, stocks paying a continuous dividend yield, options on futures, and currency options.

$$\begin{aligned}c_{GBS} &= Se^{(b-r)T}N(d_1) - Xe^{-rT}N(d_2), \\p_{GBS} &= Xe^{-rT}N(-d_2) - Se^{(b-r)T}N(-d_1),\end{aligned}\tag{4.40}$$

where

$$\begin{aligned}d_1 &= \frac{\ln(S/X) + (b + \sigma^2/2)T}{\sigma\sqrt{T}}, \\d_2 &= \frac{\ln(S/X) + (b - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}.\end{aligned}\tag{4.41}$$

and b is the cost-of-carry rate of holding the underlying security.

- $b = r$ gives the Black-Scholes (1972) stock option model,
- $b = r - q$ gives the Merton (1973) stock option model with continuous dividend yield q ,
- $b = 0$ gives the Black (1976) futures option model, and
- $b = r - r_f$ gives the Garman and Kohlhagen (1983) currency option model.

Now we will use the generalized Black-Scholes formula to evaluate examples for the different use cases presented above.

Examples: Generalized Black-Scholes Option Prices - xmpGBlackScholes

European Options on a Stock with Cash Dividends

Consider a European call option on a stock that will pay out a dividend of two, three and six months from now. The current stock price is 100, the strike is 90, the time to maturity on the option is 9 months, the risk free rate is 10% and the volatility is 25%. First calculate the stock price minus the present value of the value of the cash dividends and then use the Black-Scholes formula to calculate the call price. The result will be 15.6465.

```
# European Options on a Stock with Cash Dividends:
S <- 100 - 2*exp(-0.10*0.25) - 2*exp(-0.10*0.50)
r <- 0.10
GBlackScholes("c", S=S, X=90, Time=0.75, r=r, b=r, sigma=0.25)
[1] 15.64651
```

Options on Stock Indexes

Consider a European put option with 6 months to expiry. The stock index is 100, the strike price is 95, the risk-free interest rate is 10%, the dividend yield is 5% per annum, and the volatility is 25%. The result for the put price will be 2.4648:

```

# Options on Stock Indexes:
r <- 0.10
q <- 0.05
GBlackScholes("p", S=100, X=95, Time=0.5, r=r, b=r-q, sigma=0.20)
[1] 2.46479

```

Options on Futures

Consider a European Option on the Brent blend futures with nine months to expiry. The futures price USD 19, the risk-free interest rate is 10%, and the volatility is 28%. The result for the call price will be 1.7011 and the price for the put will be the same:

```

# Options on Futures:
FuturesPrice <- 19
b <- 0
GBlackScholes("c", S=FuturesPrice, X=19, Time=0.75, r=0.1, b=b, sigma=0.28)
1.70105
GBlackScholes("p", S=FuturesPrice, X=19, Time=0.75, r=0.1, b=b, sigma=0.28)
1.70105

```

Currency Options

Consider a European call USD put DEM option with six months to expiry. The USD/DEM exchange rate is 1.56, the strike price is 1.60, the domestic risk-free interest rate in Germany is 6%, the foreign risk-free interest rate in the United States is 8% per annum, and the volatility is 12%. The result for the call price will be 0.0291:

```

# Currency Options:
r <- 0.06
rf <- 0.08
GBlackScholes("c", S=1.56, X=1.60, Time=0.5, r=r, b=r-rf, sigma=0.12)
0.0290993

```

Our next goal is to investigate the option prices from the Black-Scholes formula as a function of their parameters. We will do this for the stock option model $b = r$.

Examples: 2D-Graphs from the Generalized Black-Scholes Option Prices - xmpGBSplots2D

Use the Splus function GBlackScholes to plot the scaled call price c/X


- as a function of S/X , the scaled asset price (note doubling the asset price and strike price yields the same option values), use as parameter $\sigma^2 T$ for a low and a high interest rate r ,
- as a function of $\sigma^2 T$, the volatility and Time to maturity (note doubling the volatility and decreasing the time to maturity by a factor of four yields the same option values), use as parameter S/X for a low and a high interest rate r .

Here is part of the Splus programming code showing how the first graph in the following figure was created:

```

S <- seq(from=0.4, to=1.2, by=0.05)
plot(S, S, ylim=c(0, 0.5), type="n", xlab="S/X", ylab="c/X")
for (Time in seq(from=1/12, to=2, by=1/12)) {
  c <- GBlackScholes("c", S=S, X=1, Time=Time, r=0.01, b=0.01, sigma=0.4)
  lines (S, c) }

```




■ Figure 4.2.1: The two upper figures show the value of a call option as a function of the underlying asset price for several times to maturity ranging from 1 to 24 months. The two lower figures show the value of a call option as a function of time to maturity for several values of the underlying asset price ranging from 0.40 to 1.20 in steps of 0.05. Call value and asset prices are measured in units of the strike. The left figures belong to low interest rate value of 1% the right figures correspond to high interest rate scenario of 10% interest rate.

Examples: 3D-Graphs from the Generalized Black-Scholes Option Prices - xmpGBSplots3D

Use the Splus function `GBSOptionPlot3D()` to plot the call price c as a function of underlying asset price S and and time to maturity $Time$ when is interest rat is 10% p.a., and the volatility is 40%. Here is the Splus command `GBSOptionPlot3D()`:

```
GBSOptionPlot3D(CallPutFlag="c", S=seq(from=75, to=125, length=40),
  X=100, Time=seq(from=1/52, to=1, length=40), r=0.1, b=0.1, sigma=0.4)
```

Here is the programming code of the `GBSOptionPlot3D()` function which shows how the plots were produced with the help of the perspective plot function `persp()` and the outer product operation function `outer()`:



■ Figure 4.2.2: The figure shows the call (left) and put (right) price for the generalized Black-Scholes model. The maturities are plotted up to 1 year for strike prices ranging from 75 to 125.

```
"GBSOptionPlot3D" <-
  function (CallPutFlag, S, X, Time, r, b, sigma) {
    "premium3D" <-
      function(S, Time, CallPutFlag, X, r, b, sigma) {
        GBSOption(CallPutFlag, S, X, Time, r, b, sigma)}
    persp(S, Time,
      outer(S, Time, FUN=premium3D, CallPutFlag, X, r, b, sigma),
      xlab="S", ylab="Time", zlab=CallPutFlag) }

```

4.2.2 Options Sensitivities

Recall from the Black-Scholes formula that the price of an option depends upon just five variables

- the current asset price,
- the strike price,
- the time to maturity,
- the volatility, and
- the interest rate.

One of these, the strike price, is normally fixed in advance and therefore does not change. That leaves the remaining four variables. We can now define four quantities, each of which measures how the value of an option will change when one of the input variables changes while the others remain the same. The definitions are as follows:

Delta

Delta means the sensitivity of the option price to the movement in the underlying asset.

$$\Delta_{call} = \frac{\partial c}{\partial S} = e^{(b-r)T} N(d_1) > 0$$

$$\Delta_{put} = \frac{\partial p}{\partial S} = e^{(b-r)T} [N(d_1) - 1] < 0$$

Delta is undoubtedly the most important measure of option price sensitivity. There are two important interpretations of Delta. One that follows directly from its definition is that delta is the slope of the premium / underlying asset price curve. A second interpretation is that Delta is the hedge ratio that can be used when hedging an option with the underlying asset. Furthermore, the Delta describes how similar the option behaves to the underlying asset. When Delta is close to zero, the option will hardly respond to movements in the underlying asset price. On the other hand when Delta approaches unity, the option moves almost one-for-one with the underlying asset, and therefore behaves very much like it.

Theta

Theta is the options sensitivity to small change in time to maturity. As time to maturity decreases, it is normal to express the Theta as minus the partial derivative with respect to time.

$$\Theta_{call} = \frac{\partial c}{\partial T} = -\frac{Se^{(b-r)T} n(d_1)\sigma}{2\sqrt{T}} - (b-r)Se^{(b-r)T} N(d_1) - rXe^{-rT} N(d_2)$$

$$\Theta_{put} = \frac{\partial p}{\partial T} = -\frac{Se^{(b-r)T} n(d_1)\sigma}{2\sqrt{T}} + (b-r)Se^{(b-r)T} N(-d_1) - rXe^{-rT} N(-d_2)$$

Thus Theta expresses how the option behaves over time. Long-dated options have more time value than short-dated ones. Therefore, as an option ages and approaches maturity, the time value will gradually erode. Theta defines exactly how much time value is lost from day to day, and is a precise measure of time decay.

Vega

The Vega is the option's sensitivity to a small movement in the volatility of the underlying asset. Note, that Vega is equal for call and put options

$$Vega_{call,put} = \frac{\partial c}{\partial \sigma} = \frac{\partial p}{\partial \sigma} = Se^{(b-r)T} n(d_1)\sqrt{T} > 0$$

Thus Vega defines the response of an option to volatility. Since higher volatility means higher uncertainty, and uncertainty manifests itself as the first component of time value, options become progressively more expensive with higher volatility.

Rho

Rho is the options sensitivity to a small change in the risk-free interest rate. For the call we have

$$\rho_{call} = \frac{\partial c}{\partial r} = TXe^{-rT} N(d_2) > 0 \text{ if } b \neq 0 ,$$

$$\rho_{call} = \frac{\partial c}{\partial r} = -Tc < 0 \text{ if } b = 0 ,$$

and for the put we have

$$\rho_{put} = \frac{\partial c}{\partial r} = -TXe^{-rT} N(-d_2) < 0 \text{ if } b \neq 0 ,$$

$$\rho_{put} = \frac{\partial c}{\partial r} = -Tp < 0 \text{ if } b = 0 .$$

Rho is probably the least used measure of sensitivity, perhaps because interest rates are relatively stable, and there is therefore less need to monitor how the option value will move when interest rates change.

All four sensitivity measures so far have one thing in common: they all express how much an option's value will change for a unit change in one of the pricing variables. Since they measure changes in premium, Delta, Theta, Vega, and Rho will all be expressed in the same units as the option premium. For example, in the case of a call option on the USD priced in DEM, the units will therefore all be fraction in DEM.

Additionally, among others, there are two other greek letters that users of options often refer to:

Lambda

Lambda or the elasticity of an option is its sensitivity in percent to a percent movement in the underlying price

$$\Lambda_{call} = \Delta_{call} \frac{S}{c} = e^{(b-r)T} N(d_1) \frac{S}{c} > 1$$

$$\Lambda_{put} = \Delta_{put} \frac{S}{p} = e^{(b-r)T} [N(d_1) - 1] \frac{S}{p} < 0$$

Thus Lambda is similar to Delta. However, instead of expressing changes in absolute terms, Lambda measures the percentage change in the premium for a percentage change in the underlying asset price. Lambda becomes thus an expression of the gearing or leverage of an option.

Gamma

Gamma is the Delta's sensitivity to small movement in the underlying asset price. Gamma is the same for call and put options.

$$\Gamma_{call,put} = \frac{\partial^2 c}{\partial S^2} = \frac{\partial^2 p}{\partial S^2} = \frac{n(d_1)e^{(b-r)T}}{S\sigma\sqrt{T}} > 0$$

Note, Gamma is the 'odd-one-out' because it is the only Greek letter that does not measure the sensitivity of the option's premium. Instead Gamma measures how the options Delta changes when the underlying asset price moves. As Delta is the single most important measure of an option's sensitivity, it makes sense to track how Delta is effected by movements in the underlying asset price. The simplest interpretation for Gamma is that it measures the curvature of the option premium graphed against the underlying asset price. Recalling that the Delta of an option is the hedge ratio, the Gamma therefore expresses how much the hedge ratio changes when the underlying asset price moves. Options with a small Gamma are therefore easy to hedge, because the hedge ratio will not change much when the underlying asset price fluctuates. Those with a high Gamma cause problems because it is constantly necessary to readjust the hedge in order to avoid risk.

The Splus library `libfOptions` implements functions to evaluate the Greeks for the generalized Black-Scholes model. We will use them in the following to calculate options sensitivities and to create plots which show the values of Greeks as function of the option parameters. First let us evaluate the Greeks for the examples given in the textbook of Haug.

Examples: Greeks for the Generalized Black-Scholes Model - xmpGBSGreeks

Delta

Consider a futures option with six months to expiry. The futures price is 105, the strike price is 100, the risk-free interest rate is 10%, and the volatility is 36%. The Delta of the call price will be 0.5946 and the Delta of the put price -0.3566 .

```
> GDelta("c", S=105, X=100, Time=0.5, r=0.1, b=0, sigma=0.36)
[1] 0.59463
> GDelta("p", S=105, X=100, Time=0.5, r=0.1, b=0, sigma=0.36)
[1] -0.35660
```


Theta

Consider an European put option on a stock index currently priced at 430. The strike price is 405, time to expiration is one month, the risk-free interest rate is 7% p.a., the dividend yield is 5% p.a., and the volatility is 20% p.a.. The Theta of the put option will be -31.1924 .

```
> GTheta("p", S=430, X=405, Time=1/12, r=0.07, b=0.07-0.05, sigma=0.2)
[1] -31.192350
```

Vega

Consider a stock option with nine months to expiry. The stock price is 55, the strike price is 60, the risk-free interest rate is 10% p.a., and the volatility is 30% p.a.. What is the Vega? The result will be 18.9358.

```
> GVega("c", S=55, X=60, Time=0.75, r=0.1, b=0.1, sigma=0.3)
[1] 18.93578
```

Rho

Consider a European call option on a stock currently priced at 72. The strike price is 75, time to expiration is one year, the risk-free interest rate is 9% p.a., and the volatility is 19% p.a.. The result for Rho will be 38.7325.

```
> GRho("c", S=72, X=75, Time=1, r=0.09, b=0.09, sigma=0.19)
[1] 38.73250
```

Lambda - Elasticity

Calculate the elasticity of a put option with the same parameters as under the Delta example. The result will be -4.8775 .

```
> GLambda("p", S=105, X=100, Time=0.5, r=0.1, b=0, sigma=0.36)
[1] -4.87751
```

Gamma

Consider the same stock option as in the example for the Vega. What will be the value of Gamma? The result for the Gamma will be 0.0278.

```
> GGamma("c", S=55, X=60, Time=0.75, r=0.10, b=0.10, sigma=0.30)
[1] 0.02782
```

The Splus functions can also be used to create 2D graphs for the sensitivities; here we display the sensitivities as function of the scaled asset price S/X , for time different times to maturity.

Examples: Greeks for the Generalized Black-Scholes Model - xmpGBSGreeks2D

As a simple illustration of how the Greeks can be used consider someone holding an at-the-money currency call on USD put on DEM with the following characteristics: underlying price 1.7000, strike price 1.7000, time to maturity 270 days, DEM (domestic) interest rate 6% p.a., USD (foreign) interest rate 3% p.a., and volatility 10% p.a.. First calculate the option value and the sensitivities for the mentioned parameters using the Splus function `GBScharacteristics`

```
GBScharacteristics(CallPutFlag="c", S=1.7000, X=1.7000,
  Time=270/365, r=0.06, b=0.06-0.03, sigma=0.10)
$premium:
[1] 0.07651851467920412
$delta:
[1] 0.6047323623055149
$theta:
[1] -0.06310373321400309
$vega:
```

```

[1] 0.5452297685828492
$rho:
[1] 0.703868918725606
$lambda:
[1] 13.43524531584737
$gamma:
[1] 2.550414783195437

```

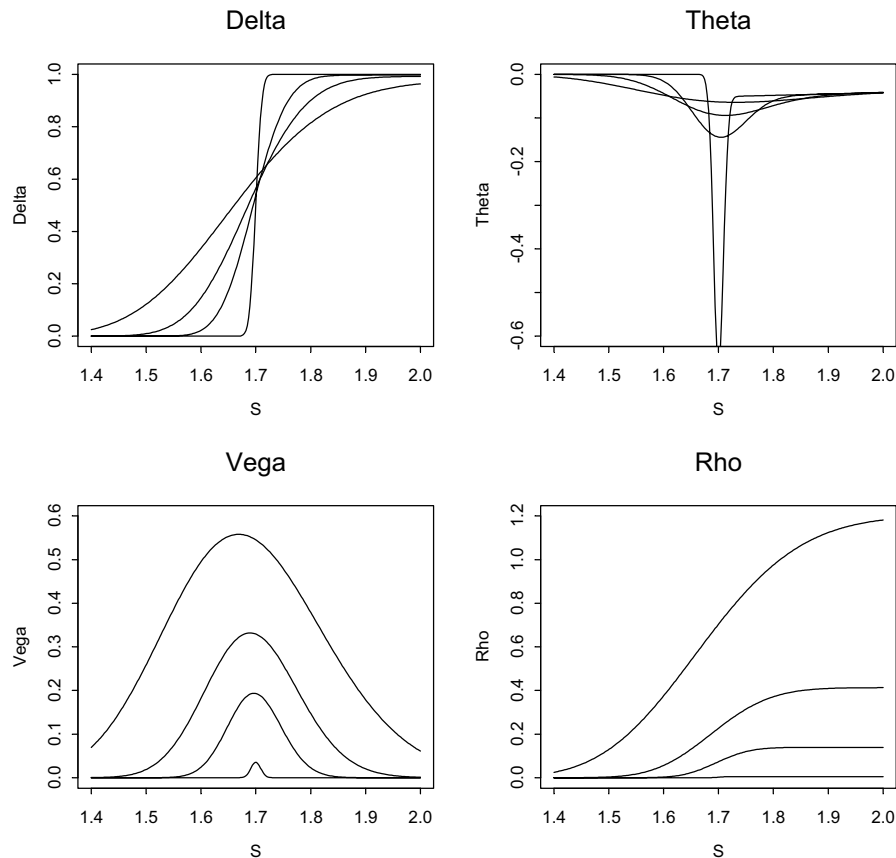
and then plot the sensitivities in the range between 1.4000 and 2.0000 for times to maturity of 1, 30, 90, and 270 days to get an overview about the whole range of prices of the underlying.

```

S <- seq(from=1.4000, to=2.0000, length=50)
Selection <- c("delta", "theta", "vega", "rho")
ymin <- c( 0.0, -0.6, 0.0, 0.0 )
ymax <- c( 1.0, 0.0, 0.6, 1.2 )
for (i in 1:4){
  plot(S, S, ylim=c(ymin[i], ymax[i]), type="n", xlab="S")
  for (Days in c(1, 30, 90, 270)) {
    c <- GGreeks(Selection[i], CallPutFlag="c", S=S, X=1.7000,
      Time=Days/365, r=0.06, b=0.06-0.03, sigma=0.1)
    lines (S, c) } }

```

The result are illustrated in figure 4.2.3.



■ Figure 4.2.3: The figures show the Delta, Theta, Vega and Gamma for four times to maturity, 1 day, 30, 90 and 270 days as a function of the strike price S .

Suppose after one week the underlying price were to rise to 1.7500 DEM, interest rates were to fall 1%, and volatility were to rise 2%. what effect would this combination have on the price of the option? Let us analyze the separate impacts, and calculate the combined effect to the option rise. The result will tell us that the effect for the premium will be to rise by 0.0329:

```

Characteristics <- GBScharacteristics(CallPutFlag = "c", S = 1.7000,
  X = 1.7000, Time = 270/365, r = 0.06, b = 0.06 - 0.03, sigma = 0.10)
Changes <- list(
  DueToPrice = (1.7500 - 1.7000) * Characteristics\$\delta,
  DueToTime = 7/365 * Characteristics\$\theta,
  DueToVolatility = (0.12 - 0.10) * Characteristics\$\vega,
  DueToInterest = (0.05 - 0.06) * Characteristics\$\rho)
print(Changes)
  $DueToPrice:
  [1] 0.03023661811527577
  $DueToTime:
  [1] -0.001210208582186361
  $DueToVolatility:
  [1] 0.01090459537165698
  $DueToInterest:
  [1] -0.007038689187256056
TotalChange <- Changes\$\DueToPrice + Changes\$\DueToTime +
  Changes\$\DueToVolatility + Changes\$\DueToInterest
print(TotalChange)
  [1] 0.03289231571749033
OldPremium <- Characteristics\$\premium
print(OldPremium)
  [1] 0.07651851467920412

```

This would imply that the premium would be $0.0765 + 0.0329 = 0.1094$. In fact if the option is repriced properly,

```

NewPremium <- GBlackScholes(CallPutFlag = "c", S = 1.7500, X = 1.7000,
  Time = 263/365, r = 0.05, b = 0.05 - 0.03, sigma = 0.12)
print(NewPremium) [1] 0.1105350734930761
RealTotalChange <- NewPremium - OldPremium
print(RealTotalChange)
  [1] 0.034016558813872

```

the premium comes out to 0.1105, an actual increase of 0.0340. This example shows, using the Greeks makes it possible to perform a quick calculation and to obtain an answer within an accuracy of about 1 to 2%. Let us finish with a the remark, that where the Greeks are becoming really important is in evaluating the impact of market fluctuations on an entire portfolio of options.

Furthermore 3d-plots can easily be created within the Splus environment allowing an even better visualization of the sensitivities in a three-dimensional asset price / time to maturity graph. This will be illustrated in the following figures.







Examples: Greeks for the Generalized Black-Scholes Model - xmpGBSGreeks3D

Plot the sensitivities Delta, Theta, Vega, Rho, Lambda and Gamma as a function of the scaled asset price S/X and the time-scaled volatility $\sigma^2 T$. Use the Splus function `GBSGreeksPlot3D`.

```

for ( Selection in c("delta", "theta", "vega", "rho", "lambda", "gamma") )
  GBSGreeksPlot3D(Selection, CallPutFlag="c",
    S=seq(from=75, to=125, length=25), X=100,
    Time=seq(from=1/52, to=1, length=25), r=0.1, b=0.1, sigma=0.40)

```



■ Figure 4.2.4: The figures show the Delta, Theta, Vega, Rho, Lambda and Gamma for the generalized Black-Scholes model as function of time to maturity and strike price.

4.2.3 Analytical Pricing Formulas for American Options

An American Option can be exercised at any time up to its expiration date. This added freedom complicates the valuation of American options relative to their European counterparts. With a few exceptions, it is not possible to find an exact formula for the value of American options. However, closed form approximations are available which are in the finance community quite popular, because the approximation formulas execute quite quickly on computers.

American Calls on Stocks with Known Dividends

Roll (1977), Geske (19979) and Whaley (1982) have developed a formula for the valuation of an American call option on a stock paying a single dividend of D , with time to dividend payout t .

$$C = (S - De^{-rt})N(b_1) + (S - Dw^{-rt})M(a_1, -b_1; -\sqrt{\frac{t}{T}}) - Xe^{-rt}M(a_2, -b_2; -\sqrt{\frac{t}{T}}) - (X - D)e^{-rt}N(b_2),$$

where

$$a_1 = \frac{\ln[(S - De^{-rt})/X] + (r + \sigma^2/2)T}{\sigma\sqrt{T}}, \quad a_2 = a_1 - \sigma\sqrt{T}$$

$$b_1 = \frac{\ln[(S - De^{-rt})/I] + (r + \sigma^2/2)T}{\sigma\sqrt{T}}, \quad b_2 = b_1 - \sigma\sqrt{T}$$

where $M(a, b; \rho)$ is the cumulative bivariate normal distribution function with upper integral limits a and b and correlation coefficient ρ . I is the critical ex-dividend stock price I that solves

$$c(I, X, T - t) = I + D - X,$$

where $c(I, X, T - t)$ is the value of the European call with stock price I and time to maturity $T - t$. If $D \leq X(1 - e^{-r(T-t)})$ or $I = \infty$, it will not be optimal to exercise the option before expiration, and the price of the American option can be found by using the European Black-Scholes formula where the stock price is replaced with the stock price minus the present value of the dividend payment $S - De^{rt}$

Example: American Calls on Stocks with Known Dividends - xmpAOPskd

Consider an American-style call option on a stock that will pay a dividend of 4 in exactly three months. The stock price is 80, the strike price is 82, time to maturity is four months, the risk-free interest rate is 6%, and the volatility is 30%. The result will be 4.3860, whereas the value of a similar European call would be 3.5107.

```
> RollGeskeWhaley(S=80, X=82, t1=1/4, T2=1/3, r=0.06, D=4, sigma=0.3)
[1] 4.38603
```

The Barone-Adesi and Whaley Approximation

The quadratic approximation method by Barone-Adesi and Whaley (1987) can be used to price American call and put options on an underlying asset with cost-of-carry rate b . When $b \geq r$, the American call value is equal to the European call value. The model is fast and accurate for most practical input values.

$$C(S, X, T) = \begin{cases} c_{GBS}(S, X, T) + A_2(S/S_1)^{q_2}, & S < S_1 \\ S - X, & S \geq S_1 \end{cases}$$

$$P(S, X, T) = \begin{cases} p_{GBS}(S, X, T) + A_1(S/S_2)^{q_1}, & S > S_2 \\ X - S, & S \leq S_2 \end{cases}$$

where c_{GBS} and p_{GBS} are the generalized Black-Scholes call and put formula, respectively, and

$$A_{1,2} = \mp \frac{S_{2,1}}{q_{1,2}} \left\{ 1 - e^{(b-r)T} N[\mp d_1(S_{2,1})] \right\},$$

$$d_1(S) = \frac{\ln(S/X) + (b + \sigma^2/2)T}{\sigma\sqrt{T}},$$

$$q_{1,2} = \frac{-(N-1) \mp \sqrt{(N-1)^2 + 4M/X}}{2},$$

$$M = 2r/\sigma^2, \quad N = 2b/\sigma^2, \quad X = 1 - e^{-rT},$$

where $S_{1,2}$ are the critical commodity prices for the call and put options respectively, that satisfy

$$S_1 - X = c(S_1, X, T) + \{1 - e^{(b-r)T} N[d_1(S_1)]\} \frac{S_1}{q_2},$$

$$X - S_2 = p(S_2, X, T) + \{1 - e^{(b-r)T} N[-d_1(S_2)]\} \frac{S_2}{q_1}.$$

Example: American Options, Barone-Adesi and Whaley Approximation - xmpAOPbaw
 Evaluate Call and Put prices of American style options using the BAW approximation formula. Choose $r=0.1$, $b=0$, and $X=100$ with volatilities σ of 15%, 25%, and 35%, and prices taking on values S of 90, 100, and 110 for two time periods, $\text{Time}=0.1$ and $\text{Time}=0.5$, respectively. The Splus function calls from the example script are:

```
# American Calls:
for ( sigma in c(0.15, 0.25, 0.35) )
  for ( S in c(90,100,110) )
    print(c(sigma, S, BAWAmericanApprox("c", S, X, Time=0.1, r, b, sigma)))
for ( sigma in c(0.15, 0.25, 0.35) )
  for ( S in c(90,100,110) )
    print(c(sigma, S, BAWAmericanApprox("c", S, X, Time=0.5, r, b, sigma)))
# American Puts:
for ( sigma in c(0.15, 0.25, 0.35) )
  for ( S in c(90,100,110) )
    print(c(sigma, S, BAWAmericanApprox("p", S, X, Time=0.1, r, 0, sigma)))
for ( sigma in c(0.15, 0.25, 0.35) )
  for ( S in c(90,100,110) )
    print(c(sigma, S, BAWAmericanApprox("p", S, X, Time=0.5, r, b, sigma)))
```

The Bjerksund and Stensland Approximation

The Bjerksund and Stensland (1993) approximation can be used to price American options on stocks, futures and currencies. The method is analytical and extremely computer efficient. Bjerksund's and Stensland's approximation is based on an exercise strategy corresponding on a flat boundary I (trigger price). Numerical investigation indicates that this model is somewhat more accurate for long-term options than the Barone-Adesi and Whaley model presented above.

$$C = \alpha S^\beta - \alpha \phi(S, T, \beta, I, I) + \phi(S, T, 1, I, I) - \phi(S, T, 1, X, I) - X \phi(S, T, 0, I, I) + X \phi(S, T, 0, X, I),$$

where

$$\alpha = (I - X)I^{-\beta}; \quad \beta = \left(\frac{1}{2} - \frac{b}{\sigma^2}\right) + \sqrt{\left(\frac{1}{2} - \frac{b}{\sigma^2}\right)^2 + 2\frac{r}{\sigma^2}}.$$

The function $\phi(S, T, \gamma, H, I)$ is given by

$$\phi(S, T, \gamma, H, I) = e^\lambda S^\gamma \left[N(d) - \left(\frac{I}{S}\right)^\kappa N\left(d - \frac{2\ln(I/S)}{\sigma\sqrt{T}}\right) \right]$$

$$\lambda = \left[-r + \gamma b + \frac{1}{2}\gamma(\gamma - 1)\sigma^2 \right] T,$$

$$d = \frac{\ln(S/H) + [b + (\gamma - \frac{1}{2})\sigma^2]T}{\sigma\sqrt{T}},$$

$$\kappa = \frac{2b}{\sigma^2} + (2\gamma - 1),$$

and the trigger price is defined as

$$I = B_0 + (B_\infty - B_0)(1 - e^{-h(T)}) ,$$

$$h(T) = -(bT + 2\sigma\sqrt{T}) \frac{\beta_0}{\beta_\infty - \beta_0} ,$$

$$B_\infty = \frac{\beta}{\beta - 1} X , \quad B_0 = \max[X, \frac{r}{r - b} X] .$$

If $S \geq I$, it is optimal to exercise the option immediately, and the value must be equal to the intrinsic value $S - X$. On the other hand, if $b \geq r$, it will never be optimal to exercise the American call option before expiration, and the value can be found using the generalized Black-Scholes formula. The value of the American put is given by the Bjerk Sund and Stensland put-call transformation

$$P(S, X, T, r, b, \sigma) = C(X, S, T, r - b, -b, \sigma) .$$

which holds in general when calculating American option prices.

Example: American Options, Bjerk Sund and Stensland Approximation - xmpAOPbsa

Consider an American-style call option with nine months to expiry. The stock price is 42, the strike price is 40, the risk free rate is 4% p.a., the dividend yield is 8% p.a., the volatility is 35% p.a.

```
> BSAmericanApprox("c", S=60, X=40, Time=0.75, r=0.04, b=0.04-0.08, sigma=0.35)
$Premium:
[1] 20
$TriggerPrice:
[1] 57.59945
```

4.2.4 Binomial Option Pricing

Beside the derivation of the Black-Scholes formula as presented in the previous sections, another approach to option valuation has become a very important tool for explaining the principles of both European and American options and also for deriving numerical solutions to option problems in practice. This is the binomial model developed by Cox, Ross and Rubinstein (1979). We will now show how this model works.

Cox-Ross-Rubinstein Binomial Tree

Consider the evaluation of an option on a non-dividend paying stock. we start by dividing the options life time into a large number of small time intervals of length Δt . We assume that in each time interval the stock price moves from its initial value of S to one of two new values, up to Su or down Sd . The probability of the up move is assumed to be p and for the down move to be $1 - p$. We apply the risk neutral valuation principle and assume that (i) the expected return from all traded securities is the risk-free interest rate, and (ii) that future cash flows can be valued by discounting their expected values at the risk free interest rate.

The parameters p , u , and d must give correct values for the mean and variance of stock price changes during a time interval Δt . The risk neutral valuation principle now requires that the expected return from a stock is the risk-free interest rate r . Hence the expected value of the stock price at the end of a time interval δt is $Se^{r\Delta t}$, where S is the stock price at the beginning of the time interval. Then it follows that

$$Se^{r\Delta t} = pSu + (1 - p)Sd . \tag{4.42}$$

The variance of the change of the stock price in a small time interval Δt is $S^2\sigma^2\Delta t$ from which it follows that

$$S^2\sigma^2\Delta t = pS^2u^2 + (1-p)S^2d^2 - S^2[pu + (1-p)d]^2 . \quad (4.43)$$

These two equations together with the condition $u = \frac{1}{d}$ which is usually used imposes three conditions on p , u , and d which imply

$$\begin{aligned} p &= \frac{e^{r\Delta t} - d}{u - d} , \\ u &= e^{\sigma\sqrt{\Delta t}} , \\ d &= e^{-\sigma\sqrt{\Delta t}} , \end{aligned} \quad (4.44)$$

provided that Δt is small.

Now let us consider the complete tree of stock prices for the binomial model. At time zero, the stock price S is known. At time Δt , there are two possible stock prices, Su and Sd ; at time $2\Delta t$ there are three possible stock prices, Su^2 , S , and Sd^2 . In general at time $i\Delta t$, $i + 1$ stock prices have to be considered. These are

$$Su^j d^{i-j} , \quad j = 0, \dots, i . \quad (4.45)$$

Then options are evaluated by starting at the end of the tree at time T and working backward. The value of the option is known at time T ; for a call $\max(S_T - X, 0)$ and for a put $\max(X - S_T, 0)$, where S_T is the stock price at time T . According to the risk-neutral valuation principle the value at each node at time $T - \Delta t$ can be calculated as the expected value at time T discounted at rate r for a time period Δt . Similarly, The value at each node at time $T - 2\Delta t$ can be calculated as the expected value at time $T - \Delta t$ discounted for a time period Δt with rate r , and so on. *Note, that if the option is American, it is necessary to check at each node to see whether early exercise is preferable to holding the option for a further time period Δt .* Eventually, by working back through all the nodes, the value of the option at time zero is obtained.

Example: Binomial Model, the Idea Behind

Consider a 5 month American put option on a non-dividend paying stock when the stock price is 50, the risk-free interest rate is 10% p.a. and the volatility is 40% p.a.. Suppose that we divide the lifetime of the option into five intervals of length one month for the purposes of constructing a binomial tree. Then $\Delta t = 0.0833$, $u = 1.1224$, $d = 0.8909$, $p = 0.5076$. The figure shows the tree.

... ..

The stock price at the j -th node ($j = 0, 1, \dots, i$) at time $i\Delta t$ is calculated as $Su^j d^{i-j}$. For example the stock price at node labeled A ($i = 4$, $j = 1$) is thus 39.69. The option prices at the final nodes are calculated as $\max(X - S_T, 0)$. For example the option price at node G is $50 - 35.36 = 14.64$. The option prices on the penultimate nodes are calculated from the option prices at the final nodes. First, we assume no exercise of the option at the nodes. This means that the option price is calculated

as the present value of the expected option price in time Δt . For example, at node E the option price is calculated as $(0.5076 \cdot 0 + 0.4924 \cdot 5.45)e^{-0.10 \cdot 0.0833} = 2.66$ while at node A it is calculated as $(0.5076 \cdot 5.45 + 0.4924 \cdot 14.64)e^{-0.10 \cdot 0.0833} = 9.90$. We then check to see if early exercise is preferable to waiting. At node E, early exercise would give a value for the option of zero, since both the stock price and strike price are 50. Clearly it is best to wait. The correct value for the option at node E is therefore 2.66. At node A it is a different story. If the option is exercised it is worth $50.00 - 39.69$ or 10.31. This is more than 9.90. If node A is reached, the option should therefore be exercised and the correct value for the option on node A is 10.31.

Option prices on earlier nodes are calculate in a similar way. Note, that it is not always best to exercise an option early when it is in-the-money. Consider node B. If the option is exercised, it is worth $50.00 - 39.69$ or 10.31. However, if it is held, it is worth $(0.5076 \cdot 6.37 + 0.4924 \cdot 14.64)e^{-0.10 \cdot 0.0833} = 10.35$. The option should therefore not be exercised at this node, and the correct option value at this node is 10.35.

Working back through the tree, we find the value of the option at the initial node to be 4.48. This is our numerical estimate for the options current value. The true value of the option, obtained using a small Δt , will become 4.29.

To find a general expression for the options value we define f_{ij} as the value of an American option at time $i\Delta t$ when the stock price is $Su^j d^{i-j}$ for $0 \leq i \leq N$, $0 \leq j \leq i$. We will refer to this as the value of the option at node (i, j) . Since the value of an American put at its expiration date is $\max(X - S_T, 0)$, we know that

$$f_{N,j} = \max[X - Su^j d^{N-j}, 0], \quad j = 0, \dots, N.$$

There is a probability, p of moving from node (i, j) at time $i\Delta t$ to the node $(i+1, j+1)$ at time $(i+1)\Delta t$, and a probability $1-p$ of moving from node (i, j) at time $i\Delta t$ to node $(i+1, j)$ at time $(i+1)\Delta t$. Assuming no early exercise, risk-neutral valuation gives

$$f_{i,j} = e^{-r\Delta t} [pf_{i+1,j+1} + (1-p)f_{i+1,j}]$$

for $0 \leq i \leq N-1$ and $0 \leq j \leq i$. When early exercise is taken into account, this value for $f_{i,j}$ must be compared with the option's intrinsic value, and we obtain

$$f_{i,j} = \max\{X - Su^j d^{i-j}, e^{-r\Delta t} [pf_{i+1,j+1} + (1-p)f_{i+1,j}]\}.$$

Note, that because the calculation start at time T and work backward, the value at time $i\Delta t$ captures not only the effect of early exercise possibilities at time $i\Delta t$, but also the effect of early exercise at subsequent times. In the limit as Δt tends to zero, an exact value for the American put is obtained. In practice $n = 30$ usually gives reasonable results.

Example: Binomial Model, the Algorithm - xmpCRRBinomial

The above formulas are implemented in the Splus function `CRRBinomial(AmeEurFlag, CallPutFlag, S, X, Time, r, b, sigma, n)` in the following form, which allows to calculate calls and puts for both, European and American options:

```
"CRRBinomial" <-
function(AmeEurFlag, CallPutFlag, S, X, Time, r, b, sigma, n) {
  if (CallPutFlag == "c") z <- +1
  if (CallPutFlag == "p") z <- -1
  dt <- Time / n
  u <- exp(sigma*sqrt(dt))
```

```

d <- 1 / u
p <- (exp(b*dt) - d) / (u - d)
Df <- exp(-r*dt)
OptionValue <- z*(S*u^(0:n)*d^(n:0) - X)
OptionValue <- (abs(OptionValue)+OptionValue)/2
if (AmeEurFlag == "e") {
  for ( j in seq(from=n-1, to=0, by=-1) )
    for ( i in 0:j )
      OptionValue[i+1] <-
        (p*OptionValue[i+2] + (1-p)*OptionValue[i+1]) * Df }
if (AmeEurFlag == "a") {
  for ( j in seq(from=n-1, to=0, by=-1) )
    for ( i in 0:j )
      OptionValue[i+1] <- max((z * (S*u^i*d^(abs(i-j)) - X)),
        (p*OptionValue[i+2] + (1-p)*OptionValue[i+1]) * Df ) }
OptionValue[1] }

```

The example script `xmpCRRBinomial()` calculates the option value for the option parameters used above for the American put and investigates the behavior in the limit $\Delta t \rightarrow 0$.

```

...
for (n in 3:30)
  OptionValue[n] <- CRRBinomial(AmeEurFlag="a", CallPutFlag="p", S=50, X=50,
    Time=0.4167, r=0.1, b=0.1, sigma=0.4, n=n)
  plot(OptionValue[3:30], type="b")
...

```

Extensions to the CRR Binomial Tree Model

The Binomial tree can be used to price options on a dividend paying stock, to price options on indices, on currencies and on futures contracts. Furthermore the approach of Cox, Ross and Rubinstein can be extended for *options on a stock paying a known dividend yield*. It can also be applied to *American barrier options* which cannot be priced analytically as their European counterparts. The binomial tree approach was also applied to *convertible bonds*. Further variants and extensions include *trinomial trees*, *three-dimensional binomial trees*, *implied binomial trees*, and *implied trinomial trees*. The Splus function Library `libfOptions` includes functions for all these cases:

```

BarrierBinomial(AmeEurFlag, TypeFlag, S, X, H, Time, r, b, sigma, n)}

ConvertibleBondBinomial(AmeEurFlag, S, X, T2, t1, r, k, q, sigma, F, Coupon, n)

Trinomial(AmeEurFlag, CallPutFlag, S, X, Time, r, b, sigma, n)

ThreeDimensionalBinomial(AmeEurFlag, TypeFlag, CallPutFlag, S1, S2, Q1, Q2, X1, X2,
  Time, r, b1, b2, sigma1, sigma2, rho, n)}

ImpliedTrinomial(ReturnFlag, STEPn, STATEi, S, X, Time, r, b, sigma, Skew, nSteps)

```

Examples are included, to present how to use these Splus functions. For a detailed description we recommend to inspect Chapters 14.2 to 14.6 in the textbook of Hull (1997) and to follow Chapters 3.1 to 3.4 in the book of Haug (1997) on option pricing formulas.

4.3 Pricing Formulas for Exotic Options

Introduction

Options are very versatile instruments and over the last years a whole range of option products has grown up, these are known as exotic options. An exotic options breaks at least one of the standard contract terms of a traditional option, notably concerning the expiry, price, strike or underlying asset(s).

An option will be valuable to the holder if it is in-the-money. How this is determined is by reference to an underlying rate. An exotic option payoff may be linked to the average underlying exchange rate over a prespecified period rather than the rate on a specific date. In contrast, the payoff may be contingent upon the price performance of a second asset, not the one on which the option is stuck. In some cases the strike may be set after the option has been set up. There are a number of ways that a traditional option structure can be manipulated to behave in a different way. Essentially, exotic options can be classified according to the following categories:

- Options with contract variations
- Path-dependent options
- Limit-dependent options
- Multi-factor options

We will discuss some selected options which are of practical use in trading. In the case of options with contract variations we present the *Binary Options* or also called Digital Options. For this kind of options the amount it is in-the-money is irrelevant, the payoff is either the predetermined amount or nothing. For the second class of path-dependent options we will describe *Asian Options* and *Lookback-Options*. With an Asian option the strike on the option is compared to an average rate over the period and with a lookback option the strike is set at maturity. Beside these simple models there exists much more complex path dependent models, known as stochastic volatility models. We will consider them later. For the third class dealing with limit-dependent options we discuss *Barrier Type Options*, which have a mechanism that activates or inactivates the option, when a particular trigger is reached. In the case of the last class concerned with multi-factor options we present *Rainbow Options* and *Quanto Options*. A rainbow call option can for example offer the holder to receive a return equal to the maximum gain from either the FTSE100, the DAX30 or the SP500 index, whereas a quanto option can for example price a foreign asset, the NIKKEI index, with a strike in USD.

We have implemented Splus functions for calculating option prices and sensitivities for the above mentioned exotic options. Now we will discuss these kind of options in some more detail.

4.3.1 Options with Contract Variations

Binary Type Options

Binary options behave similarly to standard options, but the payout is based on whether the option is on the money, not by how much it is in the money. As with a standard European style option, the payoff is based on the price of the underlying asset on the expiration date. Unlike with standard options, the payoff is fixed at the writing of the contract. Binary options may additionally incorporate barrier features. The benefits of these options are that the purchaser and writer of binary options need only determine an expected direction of price movement, rather the direction and the magnitude, in order to effectively use the option.

Use Case:

A bank wishes to hedge a key interest rate exceeding a certain level. It purchases a binary option with strike at the level at which they wish to hedge. If the interest rate exceeds that level, they receive a fixed payment, no matter how high the rate goes.

Chooser Options

The unique feature of a chooser option is the ability to purchase the option now, but not decide until later whether the option is a put or a call. Two types of chooser options exists: *Complex chooser options* with differing tenor and/or strike price for the call and the put, and *simple chooser options* with same tenor and strike for the call and the put. Chooser options are more expensive than standard options, since the purchaser has increased flexibility. Chooser options provide the benefit of allowing hedging against both price increases and decreases without purchasing both a call and a put.

Simple Chooser Option:

The simple chooser option gives the holder the right to choose whether the option is to be the standard call or put after a time t_1 , both with the same strike X and time to maturity T . The payoff from a simple chooser option at time t_1 ($t_1 < T$) is

$$w(S, X, T, t_1) = \max[c_{GBS}(S, X, T), p_{GBS}(S, X, T)] ,$$

where $c_{GBS}(S, X, T)$ and $p_{GBS}(S, X, T)$ are the Generalized Black-Scholes call and put formulas. A simple chooser option can be priced using the formula derived by Rubinstein (1991c) implemented in the Splus function

```
SimpleChooser(S, X, Time, t1, r, b, sigma)
```

Use this Splus function and calculate the price for a simple chooser option with a time to expiration of six months and time to choose between a call and a put equal to three months. The underlying stock price is 50, the strike price is 50, the risk free interest rate is 8% per annum, and the volatility per annum is 25%. The result will be $w = 6.1071$.

Complex Chooser Option:

The complex chooser option gives the holder the right to choose whether the option is to be a standard call option after a time t_1 , with time to expiration T_c and strike X_c , or put option with time to maturity T_p and strike X_p . The payoff from a complex chooser option at time t_1 ($t_1 < T_c, t_1 < T_p$) is

$$w(S, X_c, X_p, T_c, T_p, t_1) = \max[c_{GBS}(S, X_c, T_c), p_{GBS}(S, X_p, T_p)] ,$$

where $c_{GBS}(S, X, T)$ and $p_{GBS}(S, X, T)$ are the Generalized Black-Scholes call and put formulas, respectively. A complex chooser option can be priced using the formula derived by Rubinstein (1991c) implemented in the Splus function

```
ComplexChooser(S, Xc, Xp, Timec, Timep, t1, r, b, sigma)
```

Use this Splus function and calculate the price for a complex chooser option which gives the holder the right to choose whether the option is to be a call with time to expiration of six months and strike price 55, or a put with seven months to expiration and strike price 48. The time to choose between a call and a put is three months, the underlying stock price is 50, the risk free interest rate is 10% per annum, the dividend yield is 5% per annum, and the volatility per annum is 25%. The result will be $w = 6.0508$.

4.3.2 Simple Path Dependent Options

Asian Options

The unique characteristic of an Asian or average price option is that the underlying asset prices are averaged over some predefined time interval. So Asian price options are path-dependent. The price path followed by the underlying asset is crucial to the pricing of the option. The averaging tends to dampen the volatility and therefore Asian price options are less expensive than standard options. Two types of Asian price options usually are considered: (i) Asian price European options using geometric averaging, and (ii) Asian price European options using arithmetic averaging. Asian price options are commonly used in the foreign exchange, interest rate and commodity markets. There are several variations on the specifications. The averaging period can span the whole life of an option or some shorter period; options with an averaging period less than the whole life are called partial average options. The average is typically based on daily prices but could be based on weekly or monthly data. This monitoring frequency is defined in the contract. The average may be arithmetic mean (standard average), a weighted average, or a geometric mean.

Benefits: A great benefit of Asian price options is that they reduce incentives for manipulation of the underlying price at expiration. Average strike options are often used by a seller to place a floor on the selling price of a sequence of sales of an asset over some time horizon. Average strike options are cheaper than standard options. Average price options are useful in situations where the trader/hedger is concerned only about the average price of a commodity which they regularly purchase.

Use Cases:

Suppose a nine-month European average price contract calls for a payoff equal to the difference between the average price of a barrel of crude oil and a fixed exercise price of USD18. The averaging period is the last two months of the contract. The impact of this contract relative to a standard option contract is that the volatility is dampened by the averaging of the crude oil price, and

therefore the option price is lower. The holder gains protection from potential price manipulation or sudden price spikes.

A Canadian exporting firm doing business in the U.S. is exposed to CAD/USD foreign exchange risk every week. For budgeting purposes the treasurer must pick some average exchange rate in which to quote CAD cash flows (derived from USD revenue) for the current quarter. If the USD strengthens, the cash flows will be greater than estimated, but if it weakens, the company's CAD cash flows are decreased.

Geometric Average-Rate Options:

If the underlying asset is assumed to be lognormally distributed, the geometric average $(x_1 + \dots + x_n)^{1/n}$ of the asset will itself be lognormally distributed. As originally shown by Kemna and Vorst (1990), the geometric average option can be priced as a standard option by changing the volatility $\sigma_A = \sigma/\sqrt{3}$ and cost of carry term $b_A = \frac{1}{2}(b - \frac{\sigma^2}{6})$.

A geometric average-rate option can be priced using the Splus function

```
GeometricAverageRateOption(CallPutFlag, S, SA, X, Time, Time2, r, b, sigma)
```

What is the value of a geometric average-rate put option with three months to maturity? The strike is 85, the asset price is 80, the risk free rate is 5%, the cost of carry is 8%, and the volatility is 20%. The result will be $p = 4.6922$. Note, that the value of the standard European put option is $p = 5.2186$.

Arithmetic Average-Rate Options:

It is not possible to find a closed form solution for the valuation of options on an arithmetic average $(x_1 + \dots + x_n)/n$. The main reason for this is that when the asset is assumed to be lognormally distributed, the arithmetic average will not itself have a lognormal distribution. Arithmetic average-rate options can be priced by analytical approximations or by numerical methods.

Turnbull's and Wakeman's Approximation: This approximation (1991) adjusts the mean, $b_A = \frac{\ln M_1}{T}$ and variance $\sigma_A^2 = \frac{\ln M_2}{T} - 2b_A$ so that they are consistent with the exact moments M_1 and M_2 of the arithmetic average, some lengthy expressions.

Levy's Approximation: An alternative formula yields this approximation (1992) with approximative quantities for S, X, d_1 and d_2 . Although the formula does not allow for $b = 0$, Levy's formula is expected to be a bit more accurate compared to the first approximation of Turnbull and Wakeman.

Curran's Approximation: This approximation (1992) is based on a method called geometric conditioning approach, claiming to be more accurate than other closed-form approximations derived earlier.

Use the following three Splus functions and compare the three approximations in a reasonable range for the parameters

```
TurnbullWakemanAsian(CallPutFlag, S, SA, X, Time, Time2, tau, r, b, sigma)
LevyAsian(CallPutFlag, S, SA, X, Time, Time2, r, b, sigma)
```

Lookback Options

The lookback option is unique because it gives the holder the right to buy an asset at its lowest price or sell it at its highest price attained over the life of the option. At expiration you "look back" and choose the best price that occurred during the option term. For a lookback call option, the lowest observed price is selected and is applied as the strike exercise price. For a lookback put option, the highest price is selected and is applied as the exercise price. The holder

of a lookback option can never miss the best underlying asset price. These options reduce regret, since they guarantee a payoff if the option is in-the-money at any point during its life.

Benefits: These options pay the largest in-the-money amount over the life of the option. The lookback call owner can buy at the lowest observed price or rate. The lookback put owner can sell at the highest observed price or rate. A lookback option can never be out-of-the money. The lookback holder gains economic value through hindsight.

Features: The lookback and standard call option prices converge as the underlying price increases. A lookback option is more expensive than a standard option. A justification for this higher cost is that they optimize market timing- a call option provides the best timing when awaiting an increase in the underlying asset price while the put gives the best timing when expecting a downturn. An in-the-money lookback approaches the value of the standard option.

Use Cases:

Consider a six-month currency lookback call option on 1 million BP against US dollars. At option expiration, you can lookback over the preceding six months and chose to accept sterling at the most favorable exchange rate that occurred. This guarantees the a no-regrets result since the best exchange rate will be achieved.

A US manufacturer buys raw materials from a Canadian supplier. Upon receipt, he has until months end to settle and is thus exposed to foreign exchange risk on a monthly basis. The manufacturer would like to lock-in the most favorable exchange rate in that monthly interval.

Currency-linked bond issues, to avoid missing the best currency rates.

Open-end offshore investment funds, to assure each new investor a lock-in of the best currency rates throughout participation in the fund.

Floating Strike Lookback Options:

A floating strike lookback call gives the holder of the option the right to buy the underlying security at the lowest price observed, min , in the life of the option. Similarly, a floating strike lookback put gives the option holder the right to sell the underlying security at the highest price observed, S_{max} in the option's life time. The payoff from a floating strike lookback call and put options are

$$c(S, S_{min}, T) = \max(S - S_{min}; 0) = S - S_{min} ,$$

$$c(S, S_{max}, T) = \max(S_{max} - S; 0) = S_{max} - S .$$

Floating strike lookback options were originally introduced by Goldman, Sosin and Gatto (1997). Their result for the option prices is implemented in the Splus function

```
FloatingStrikeLookback(CallPutFlag, S, X, T, r, b, sigma)
```

Fixed Strike Lookback Options:

In a fixed strike lookback call, the strike is fixed in advance, and at expiry the the option pays out the maximum of the difference between the highest observed price, S_{max} , in the option lifetime and the strike X , and 0. Similarly, a put at expiry pays out the maximum of the difference between the fixed strike X and the minimum observed price, S_{min} , and 0. Fixed price lookback options can be priced using the Conze and Viswanathan (1991) formula, which is implemented in the Splus function

```
FixedStrikeLookback(CallPutFlag, S, X, T, r, b, sigma)
```

Partial-Time Floating Strike Lookback Options:

In the partial-time floating strike lookback options, the lookback period is at the beginning of the options's lifetime. Time to expiry is T_2 , and time to the end of the lookback period is t_1 ($t_1 \leq T_2$). Except for the partial lookback period, the partial-time floating strike lookback option is similar to a standard floating strike lookback option. However, a partial lookback option must naturally be cheaper than a similar standard floating strike lookback option. Heynen and Kat (1994) have developed formulas for pricing partial-time floating strike lookback options which are implemented in the Splus function

`PartialFloatLB(CallPutFlag, S, X, T, r, b, sigma)`

Partial-Time Fixed Strike Lookback Options:

In the partial-time fixed strike lookback options, the lookback period starts at a predetermined time t_1 after the initialization time of the option. The partial-time fixed strike lookback call gives a payoff equal to the maximum of the highest observed price of the underlying asset, S_{max} in the lookback period minus the strike price X , and 0. The put gives a payoff equal to the maximum of the fixed strike price X minus the minimum observed asset price, S_{min} in the lookback period, $T_2 - t_1$, and 0. The partial-time fixed strike lookback option is naturally cheaper than a similar standard fixed strike lookback option. Heynen and Kat (1994) have developed formulas for pricing partial-time fixed strike lookback options which are implemented in the Splus function

`PartialFixedLB(CallPutFlag, S, X, T, r, b, sigma)`

4.3.3 Limit Dependent Options

Barrier Options

Barrier options are similar to standard options except that they are extinguished or activated when the underlying asset price reaches a predetermined barrier or boundary price. Barrier options are also path-dependent since they are dependent on the price movement of the underlying asset. A knock-out option will expire early if the barrier price is reached whereas a knock-in option will come into existence if the barrier price is reached. As with average options, a monitoring frequency is defined as part of the option which specifies how often the price is checked for breach of the barrier. The frequency is normally continuous but could be hourly, daily, etc.

The following situations can be distinguished:

- *Down and Out*: The option is canceled or knocked-out if the asset falls to a predetermined boundary price.
- *Down and In*: The option is activated or knocked-in if the asset falls to a predetermined boundary price.
- *Up and Out*: The option is canceled or knocked-out if the asset rises to a predetermined boundary price.
- *Up and In*: The option is activated or knocked-in if the asset rises to a predetermined boundary price.

Benefits: The premium for barrier options is lower than standard options, as the barrier option will have value within a smaller price range than the standard option. The owner of a barrier

option loses some of the traditional option value and therefore it should sell at a lower price than a standard option. The seller of the barrier lowers his exposure or risk, relative to a standard option. Some barrier options offer a rebate; should the option be knocked-out, the holder would receive a predefined payoff. This feature is less common. Obviously a barrier option with rebate has more value than one without.

Features: What do you have when you buy a down-and-out call and a down-and-in call, assuming no rebate? A standard call! Barrier options are hybrids: they are European in that they could have a payoff at expiration but they are American in that they may be exercised (extinguished) prior to expiration. The lower cost of the barrier option makes it one of the most popular of the exotic options for hedging purposes. Speculators are able to gain greater leverage with barrier options for the same dollar amount.

Use Cases:

A bank may wish to purchase an at-the-money 9-month Nikkei call option struck at 17,000 with a down-and-out barrier price of 16,000. If the price of the Nikkei falls to 16,000 or below, during the 9-month period, the bank will no longer have the benefit of Nikkei price appreciation since the call option will have been knocked out.

An airline is concerned that events in the Middle East might drive up the price of fuel. An up-and-in call would allow the airline to buy crude oil futures at a fixed price if some knock-in boundary price is reached. The price of the up-and-in call would be less than a standard call with the same expiration and exercise price so it might be viewed as a cost effective hedging instrument.

4.3.4 Multiple Factors Options

Rainbow Options

The term rainbow option is applied to an entire class of options which are written on more than one underlying asset. Rainbow options are usually calls or puts on the best or worst of n underlying assets, or options which pay the best or worst of n assets. Spread options are a special case of rainbow options.

Benefits: Rainbow options are tools for hedging the risk of multiple assets. Rainbow options provide effective hedging of assets with negative correlation. A put on the worse of two assets provides protection from price movements in either direction.

Features: Options on two highly correlated assets are less expensive than options on two assets which are not correlated, as lower correlation implies more variability in the individual prices. Rainbow options at exercise may deliver either the best or worse asset in the rainbow or a call or put option on the better or worse of the assets. Multi-color rainbow options could deliver the best or worst m of the n assets. Rainbow model will price two-color rainbow options (options on two underliers). Options on more than two underliers are less common and in general not analytically tractable.

Use case:

An investment manager holds two negatively correlated risky assets. The manager expects a price movement but does not know which asset will increase in price and which will decrease. A put on the worse of the two assets will provide insurance no matter which price movement occurs.

4.3.5 Pricing Formulas for Other Exotic Options

This was only a small number of exotic option models. In the book *The Complete Guide to Option Pricing Formulas* E.G. Haug (1979) presents a large collection of option pricing formulas. We have implemented all the different exotic option pricing formulas from the models presented in his book into Splus functions. A work currently under progress is to implement also the option sensitivities for all the models.

Here is a brief list of all Splus functions to price exotic options.

```
Executive(CallPutFlag, S, X, T, r, b, sigma)
ForwardStartOption(CallPutFlag, S, X, T, r, b, sigma)
TimeSwitchOption(CallPutFlag, S, X, T, r, b, sigma)
SimpleChooserOption(CallPutFlag, S, X, T, r, b, sigma)
ComplexChooserOption(CallPutFlag, S, X, T, r, b, sigma)
OptionsOnOptions(CallPutFlag, S, X, T, r, b, sigma)
ExtendibleWriter(CallPutFlag, S, X, T, r, b, sigma)
TwoAssetCorrelation(CallPutFlag, S, X, T, r, b, sigma)
EuropeanExchangeOption(CallPutFlag, S, X, T, r, b, sigma)
AmericanExchangeOption(CallPutFlag, S, X, T, r, b, sigma)
ExchangeExchangeOption(CallPutFlag, S, X, T, r, b, sigma)
OptionsOnTheMaxMin(CallPutFlag, S, X, T, r, b, sigma)
SpreadApproximation(CallPutFlag, S, X, T, r, b, sigma)
FloatingStrikeLookback(CallPutFlag, S, X, T, r, b, sigma)
FixedStrikeLookback(CallPutFlag, S, X, T, r, b, sigma)
PartialFloatLB(CallPutFlag, S, X, T, r, b, sigma)
PartialFixedLB(CallPutFlag, S, X, T, r, b, sigma)
ExtremeSpreadOption(CallPutFlag, S, X, T, r, b, sigma)
StandardBarrier(CallPutFlag, S, X, T, r, b, sigma)
DoubleBarrier(CallPutFlag, S, X, T, r, b, sigma)
PartialTimeBarrier(CallPutFlag, S, X, T, r, b, sigma)
TwoAssetBarrier(CallPutFlag, S, X, T, r, b, sigma)
PartialTimeTwoAssetBarrier(CallPutFlag, S, X, T, r, b, sigma)
LookBarrier(CallPutFlag, S, X, T, r, b, sigma)
SoftBarrier(CallPutFlag, S, X, T, r, b, sigma)
GapOption(CallPutFlag, S, X, T, r, b, sigma)
CashOrNothing(CallPutFlag, S, X, T, r, b, sigma)
TwoAssetCashOrNothing(CallPutFlag, S, X, T, r, b, sigma)
AssetOrNothing(CallPutFlag, S, X, T, r, b, sigma)
SuperShare(CallPutFlag, S, X, T, r, b, sigma)
BinaryBarrier(CallPutFlag, S, X, T, r, b, sigma)
GeometricAverageRateOption(CallPutFlag, S, X, T, r, b, sigma)
TurnbullWakemanAsian(CallPutFlag, S, X, T, r, b, sigma)
LevyAsian(CallPutFlag, S, X, T, r, b, sigma)
FourEquOptInDomCur(CallPutFlag, S, X, T, r, b, sigma)
Quanto(CallPutFlag, S, X, T, r, b, sigma)
EquityLinkedFXO(CallPutFlag, S, X, T, r, b, sigma)
TakeoverFXoption(CallPutFlag, S, X, T, r, b, sigma)
```

4.4 Heston Nandi Option Pricing

Introduction

The classical Black-Scholes model is the most popular model for the valuation and the risk-management of derivative securities. In the classical Black-Scholes model the prices of an asset are assumed to be paths of a geometric Brownian motion process. The price c_{BS} of an European call option with this asset as its underlying is then evaluated as:

$$c_{BS}(S, \sigma, t) = S N(d_t^1) - X e^{-r(T_m-t)} N(d_t^2) , \quad (4.46)$$

where

$$d_1 = \frac{\ln(S/X) + (r + \sigma^2/2)(T_m - t)}{\sqrt{(T_m - t)}\sigma} , \quad d_2 = d_1 - \sqrt{(T_m - t)}\sigma . \quad (4.47)$$

S denotes the current asset price, X the strike price of the call option, $T_m - t$ is the time to maturity T^1 , and r is the risk free interest rate. The volatility σ is the standard deviation of the log returns of the asset prices. N denotes the cumulative distribution function for a standard normal variable.

One of the drawbacks of the classical Black-Scholes model is the assumption of a constant volatility of the log return time series of the underlying asset price. In recent years there have been many approaches to include a time varying and even a stochastic volatility of the asset price time series into the concept of option pricing. The most popular class of stochastic processes for the modelling of the financial market time series are the many variants of GARCH processes. Unfortunately many of these models lack a simple option pricing formula, in most cases the price of an option even has to be computed with the help of numerical simulations.

However, Heston and Nandi (1999) formulated a special kind of GARCH process for which a closed formula for the prices of European call options exists. The Heston-Nandi GARCH model exhibits several effects known from the real financial markets, which cannot be explained by the Black-Scholes model:

¹In the following sometimes the index m may be suppressed.

- *Leptokurtosis* The (unconditional) distribution of the asset price log returns is heavy tailed.
- *Memory*: The autocorrelations of the squared and absolute log returns of the spot prices decay much slower compared with an exponential rate.
- *Clustering of volatility*: There are periods where the volatility is high, followed by periods where the volatility is low. Especially large or especially small conditional variances persist over a certain time interval.
- *Leverage Effect*: The conditional variance and the asset price log returns are negatively correlated.
- *Smile effect*: The Black-Scholes implied volatilities observed for real financial market data show a characteristic "smile" pattern. Additionally, the *skew* of the smile pattern can be produced by taking the leverage effect into consideration.

Therefore the Heston-Nandi GARCH option pricing model might be used to study and to model more closely the behavior of real financial market data. Moreover, it can be used to test numerical approaches, such as Monte Carlo simulations, applied to path dependent options.

First we formulate the Heston-Nandi GARCH(1,1) process and present the associated formula for the pricing of European call options. Moreover, we state the expressions for the Greeks of Heston-Nandi call options and introduce a hedge ratio for the hedging of derivatives within the Heston-Nandi framework. Then we present the functions of the Splus library for the Heston-Nandi analysis.

4.4.1 The Heston Nandi GARCH Option Pricing Model

In the Heston-Nandi GARCH model the asset prices are modelled by the discrete time stochastic process S_t for which the associated log return process

$$Y_t = \log \frac{S_t}{S_{t-1}} = \log S_t - \log S_{t-1} \quad (4.48)$$

define a so called Heston-Nandi GARCH process. In the following we will focus on the Heston Nandi GARCH(1,1) model.

The Heston-Nandi GARCH(1,1) Process

The Heston-Nandi GARCH(1,1) process is defined by

$$Y_t = r + \lambda \sigma_t^2 + \sigma_t Z_t, \quad (4.49)$$

$$\sigma_t^2 = \omega + \alpha (Z_{t-1} - \gamma \sigma_{t-1})^2 + \beta \sigma_{t-1}^2, \quad (4.50)$$

with $\omega > 0$, $\alpha > 0$ and $\beta > 0$. The *innovations* Z_t are *iid.* random variables with $E[Z_t] = 0$ and $E[Z_t^2] = 1$; in the following we assume the innovations to be standard normally distributed. The *conditional variances* σ_t^2 of the log returns between two discrete time steps are itself modelled by a stochastic process. The conditional variance σ_t^2 during the time interval $[t-1, t)$ depends on the information set at time $t-1$ but is independent on the innovation Z_t at time t . r is

the *continuously compounded interest rate* for the time interval between two successive discrete times. Additionally the conditional variance appears in the mean as a return premium, which allows the conditional mean of the log return to depend on the level of risk.

In contrast to other GARCH processes the conditional variance process is driven directly by the innovations Z_t and not by the log returns $\sigma_t Z_t$. This may be considered as a simplification of the ordinary GARCH processes, which allows for the existence of a closed option pricing formula. When the parameters α and β are set zero, the HN GARCH(1,1) process coincides with the discrete time geometric Brownian motion process, which is the stochastic process for the asset prices in the discrete time Black-Scholes model.

The HN GARCH(1,1) process is stationary if $\beta + \alpha\gamma^2 < 1$. The *persistence* $\beta + \alpha\gamma^2$ is essentially the quantity which measures the correlation between two successive absolute or squared log returns, or in other words the memory of the absolute or squared HN GARCH(1,1) time series. For a stationary HN GARCH(1,1) process the *unconditional variance* σ^2 , which is the expectation value of the conditional variance can be evaluated as

$$E[\sigma_t^2] = \frac{\alpha + \omega}{1 - \beta - \alpha\gamma^2} . \quad (4.51)$$

The unconditional variance corresponds to the squared volatility in the Black-Scholes model. The parameter γ is responsible for the (conditional) correlation between the log returns and the conditional variances

$$\text{Cov}_{t-1}[\sigma_{t+1}^2, Y_t] = -2\alpha\gamma\sigma_t^2 . \quad (4.52)$$

Therefore the parameter γ is the appropriate quantity for the modelling of the leverage effect. For the stationary HN GARCH(1,1) processes the heavy tails of the (unconditional) log return distribution can be measured by the *kurtosis* of the log returns

$$k = \frac{E[(Y - E[Y])^4]}{E[(Y - E[Y])^2]^2} . \quad (4.53)$$

The Heston-Nandi GARCH(1,1) European Call Option Pricing Formula

Following Brennan (1979) and Rubinstein (1976), it can be shown that in the Heston Nandi GARCH(1,1) framework the following holds: If the price process of the underlying asset is assumed to be risk neutral, the price of a derivative at time t is equal to the expectation value of the payoff of the derivative at maturity T , conditioned on the events up to time t and discounted at the risk free interest rate r . If the price process of the underlying is not risk-neutral, the risk neutral price process has to be used for the computation of the price of an option.

In order to calculate the price of a derivative the *conditional generating function* of the spot price process

$$f(\Phi) = E_t[S_T^\Phi] . \quad (4.54)$$

has to be computed in a first step. For the HN GARCH(1,1) process the conditional generating function assumes the following form

$$f(\Phi) = S_T^\Phi \exp \left(A(t; T, \Phi) + B(t; T, \Phi) \sigma_{t+1}^2 \right) , \quad (4.55)$$

where

$$\begin{aligned} A(t; T, \Phi) &= A(t+1, T, \Phi) + \Phi r + B(t+1; T, \Phi) \omega - \frac{1}{2} \ln(1 - 2\alpha B(t+1; T, \Phi)) , \\ B(t; T, \Phi) &= \Phi (\lambda + \gamma) - \frac{1}{2} \gamma^2 + \beta B(t+1; T, \Phi) + \frac{\frac{1}{2}(\Phi - \gamma)^2}{1 - 2\alpha B(t+1; T, \Phi)} , \end{aligned} \quad (4.56)$$

which can be computed recursively from the boundary conditions

$$A(t-1; T, \Phi) = \Phi r , \quad (4.57)$$

$$B(t-1; T, \Phi) = \lambda \Phi + \frac{1}{2} \Phi^2 . \quad (4.58)$$

Since $f(i\Phi)$ is the characteristic function of the log spot price it is possible to calculate the requested expectation values by inverting the characteristic function following Feller (1971) [?] or Kendall and Stuart (1977) [?]. Since the price of a derivative is the expectation value of the payoff at maturity under the risk neutral process, the conditional generating function has to be calculated for the risk neutral process. For a HN GARCH(1,1) process as defined in equations (4.49) and (4.50) the associated risk neutral process assumes the same form only with the following parameters replaced by:

$$\lambda^* = -\frac{1}{2} , \quad (4.59)$$

$$\gamma^* = \gamma + \lambda + \frac{1}{2} . \quad (4.60)$$

In the second step the Heston-Nandi European call option pricing formula for a discrete time HN GARCH(1,1) process can be evaluated: At time t , an European call option with strike price X and expiring at maturity time T is worth

$$c = e^{-r(T-t)} E_t^* [\max(S_T - X, 0)] = \quad (4.61)$$

$$\begin{aligned} & \frac{1}{2} S_t + \frac{e^{-r(T-t)}}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} f^*(i\Phi + 1)}{i\Phi} \right] d\Phi \\ & - X e^{-r(T-t)} \left(\frac{1}{2} + \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} f^*(i\Phi)}{i\Phi} \right] d\Phi \right) , \end{aligned}$$

where $E_t^*[\cdot]$ denotes the expectation value and $f^*(\Phi)$ the conditional generating function under the risk-neutral process.

The Greeks of a Heston-Nandi GARCH(1,1) European Call Option

The HN call option price is a function of the asset price S , the strike price X , the time to maturity $(T - t)$, the risk free interest rate r , and the conditional variance σ_{t+1}^2 during the time interval $[t, t + 1)$. Similar to the expression for the price of an European call option, there exist closed form expressions for the Greeks, i.e. expressions for the several partial derivatives of the Heston-Nandi call option price with respect to these variables.

The delta of an European call option can be computed straightforwardly by differentiating the expression (4.61) for the call option price with respect to the asset price S . Assuming that it is allowed, to interchange the differentiation and the integration in equation (4.61), the integrals for the delta assume a similar form as in the formula for the call option price, only that the conditional generating function $f^*(\Phi)$ for the risk neutral process is replaced by its partial derivative with respect to the asset price $\frac{\partial f^*}{\partial S}$:

$$\Delta = \frac{\partial c}{\partial S} = \quad (4.62)$$

$$\frac{1}{2} + \frac{e^{-r(T-t)}}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} \frac{\partial f^*(i\Phi+1)}{\partial S}}{i\Phi} \right] d\Phi - X e^{-r(T-t)} \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} \frac{\partial f^*(i\Phi)}{\partial S}}{i\Phi} \right] d\Phi ,$$

with

$$\frac{\partial f^*}{\partial S}(\Phi) = \frac{\Phi}{S} f^*(\Phi) .$$

Similarly we compute the expression for the gamma of the HN call option. This time the conditional generating function in the integral for the gamma has to be replaced by $\frac{\partial^2 f^*}{\partial S^2}$:

$$\Gamma = \frac{\partial^2 c}{\partial S^2} = \quad (4.63)$$

$$\frac{e^{-r(T-t)}}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} \frac{\partial^2 f^*(i\Phi+1)}{\partial S^2}}{i\Phi} \right] d\Phi - X e^{-r(T-t)} \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} \frac{\partial^2 f^*(i\Phi)}{\partial S^2}}{i\Phi} \right] d\Phi ,$$

with

$$\frac{\partial^2 f^*}{\partial S^2}(\Phi) = \frac{\Phi(\Phi - 1)}{S^2} f^*(\Phi) .$$

The expression for the partial derivative of the call option price with respect to the risk free interest rate r , the rho, is obtained by differentiating the expression (4.61) for the call option price with respect to r :

$$\begin{aligned} \rho = \frac{\partial c}{\partial r} = & -(T-t) \frac{e^{-r(T-t)}}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} (1-\Phi) f^*(i\Phi+1)}{i\Phi} \right] d\Phi \\ & + (T-t) X e^{-r(T-t)} \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} (1-\Phi) f^*(i\Phi)}{i\Phi} \right] d\Phi . \end{aligned} \quad (4.64)$$

Since the HN GARCH(1,1) model, discussed here, is a discrete time model, there exists no derivative of the call option price with respect to the time t . However, the theta can be approximated, by taking the differences between the call option prices, that are separated by one time step and then dividing by the length of the corresponding time interval, which is in our case just one time unit.

In the case where the structure parameters α , β and γ are set zero, the HN GARCH(1,1) model and the BS model coincide, if the squared volatility in the BS model is equal to the unconditional variance in the HN GARCH(1,1) model, that is $\sigma^2 = \omega$. Then the delta, the gamma, the rho and the theta are identical in both models. However, the vega of the two models do not coincide. The vega in the BS model is the partial derivative of the call option price with respect to the volatility σ . The corresponding quantity in the HN GARCH(1,1) model were the partial derivative of the HN call option price with respect to the square root of the unconditional variance. In contrast to this, the vega in the HN GARCH(1,1) model is the partial derivative with respect to the conditional variance σ_{t+1}^2 during the time interval $[t, t+1)$:

$$\text{Vega} = \frac{\partial c}{\partial \sigma_{t+1}^2} = \quad (4.65)$$

$$\frac{e^{-r(T-t)}}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} \frac{\partial f^*(i\Phi+1)}{\partial \sigma_{t+1}^2}}{i\Phi} \right] d\Phi - X e^{-r(T-t)} \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{X^{-i\Phi} \frac{\partial f^*(i\Phi)}{\partial \sigma_{t+1}^2}}{i\Phi} \right] d\Phi ,$$

with

$$\frac{\partial f^*}{\partial \sigma_{t+1}^2}(\Phi) = B(t; T, \Phi) f^*(\Phi) .$$

The partial derivatives with respect to the structure parameters $\{\lambda, \alpha, \beta, \omega, \gamma\}$ can be found similarly. Let Ω be one of the structure parameter under consideration. Then the following holds:

$$\frac{\partial c}{\partial \Omega} = \quad (4.66)$$

$$\frac{e^{-r(T-t)}}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{K^{-i\Phi} \frac{\partial f^*(i\Phi+1)}{\partial \Omega}}{i\Phi} \right] d\Phi - K e^{-r(T-t)} \frac{1}{\pi} \int_0^\infty \operatorname{Re} \left[\frac{K^{-i\Phi} \frac{\partial f^*(i\Phi)}{\partial \Omega}}{i\Phi} \right] d\Phi ,$$

where

$$\frac{\partial f^*}{\partial \Omega}(\Phi) = \left(\frac{\partial A}{\partial \Omega}(t; T, \Phi) + \sigma_{t+1}^2 \frac{\partial B}{\partial \Omega}(t; T, \Phi) \right) f^*(\Phi) .$$

Again $\frac{\partial A}{\partial x}(t; T, \Phi)$ and $\frac{\partial B}{\partial x}(t; T, \Phi)$ can be computed by recursion relations, which simply can be found by differentiating the initial recursion relation (4.56), with respect to the corresponding parameter.

Since the price of a derivative is always computed under the risk neutral measure, the partial derivative of the call option price with respect to the structure parameter λ is equal to the derivative with respect to the parameter γ . This follows from from equation (4.60).

Hedging in the Heston-Nandi GARCH(1,1) Model

The delta hedging method which is suitable for the Black-Scholes model is not the best hedging strategy in the HN GARCH(1,1) model. This is because the HN GARCH(1,1) model is a stochastic volatility model and therefore has additional variability in the asset price process. Note that the risk of a portfolio can never be eliminated perfectly, since the HN GARCH(1,1) model is a discrete time model.

A suitable *hedge ratio* h for the HN GARCH(1,1) model can be found by minimizing the variance of the *hedge error*

$$\Delta \Pi = -\Delta f + h \Delta S , \quad (4.67)$$

where Δf denotes the change of the price of a derivative and ΔS denotes the change of the asset price during one time step. The minimum variance hedge ratio can be approximated (with a first order Taylor expansion of the price of the derivative) as

$$h_t^{Var} = \frac{\partial f}{\partial S} - \frac{2\alpha\gamma}{S_t} \frac{\partial f}{\partial \sigma_{t+1}^2} + \frac{2\alpha^2}{S_t \sigma_{t+1}^2 \left(\frac{\partial f}{\partial S} S_t - 2\alpha\gamma \frac{\partial f}{\partial \sigma_{t+1}^2} \right)} \frac{\partial f}{\partial \sigma_{t+1}^2} . \quad (4.68)$$

Another suitable hedge ratio can be gained by eliminating the contribution in $\Delta \Pi$ which depends linearly on the change ΔS of the asset price. A straightforward calculation yields for this hedge ratio

$$h_t = \frac{\partial f}{\partial S}(S_t, \sigma_{t+1}^2, t+1) - \frac{2\alpha}{S_t} \left(\gamma + \lambda + \frac{\mu}{\sigma_{t+1}^2} \right) \frac{\partial f}{\partial \sigma_{t+1}^2}(S_t, \sigma_{t+1}^2, t+1) . \quad (4.69)$$

In the following we use this expression as the hedge ratio² in the HN GARCH(1,1) model and refer to it as the HN hedge ratio.

²The approximation (4.68) for the minimum variance hedge ratio is numerically unstable, because the first order Taylor expansion is not enough accurate.

4.4.2 Numerical Analysis of the Heston Nandi Model

The Splus function package for the Heston-Nandi analysis of financial market time series includes functions for the simulation of HN GARCH(1,1) asset price paths, the estimation of the structure parameters of a HN GARCH(1,1) process, the calculation of the first four moments of the (unconditional) log return distribution of a given HN GARCH(1,1) process, the recalculation of the conditional variances and the associated innovations from a time series, the pricing and hedging of European call options according to Heston-Nandi and according to Black-Scholes, and the extraction of BS implied volatilities. In the following we present the functions included in the `fOptions` library, and demonstrate its applications at hand of several examples.

Overview over the Software Package

Before discussing the Splus functions in detail, we give a short overview over the functions included in the `fOptions` library:

Heston Nandi GARCH(1,1) functions

- `hngarch.mle` - max log likelihood estimation of the HN GARCH(1,1) structure parameters from observed log return time series
- `hngarch.sim` - simulation of HN GARCH(1,1) time series
- `hngarch.diag` - recalculation of the conditional variances and the associated innovations of a HN GARCH(1,1) process from an observed log return time series
- `hngarch.mom` - moments of the HN GARCH(1,1) (unconditional) log return distribution
- `hngarch.momdiff` - partial derivatives of the moments of the HN GARCH(1,1) (unconditional) log returns with respect to the HN GARCH(1,1) structure parameters
- `hngarch.rms` - root mean square error estimates of the HN GARCH(1,1) structure parameters from observed options prices

- `HNGOption` - call and put prices of HN GARCH(1,1) options
- `HNGGreeks` - the Greeks of the HN GARCH(1,1) call and put options
- `HNGOmegas` - partial derivatives of the HN GARCH(1,1) call and put option prices with respect to the structure parameters $\lambda, \alpha, \beta, \omega$ and γ

In analogy we have also added the Black-Scholes functions

- `bs.imv` - extraction of BS implied volatilities from the call option prices
- `bs.rms` - root mean squared error estimate of BS volatility from observed call option prices

The notations `hngarch.*` follows those from standard Splus time series analysis and the notations `HNG*` follow those from the Black-Scholes options pricing.

Maximum Log Likelihood Parameter Estimation

The Splus function `hngarch.mle(model, x, symmetric=T, doprint=T)` allows to estimate the structure parameters $(\lambda, \omega, \alpha, \beta, \gamma)$ of a HN GARCH(1,1) process from an observed financial

market log return time series. The parameters are estimated via the maximum loglikelihood approach, maximizing the log-likelihood function of the HN GARCH(1,1) process, assuming that the innovations are Gaussian distributed. `hngarch.mle()` returns the list `$model=list(lambda, omega, alfa, beta, gamma)`, which contains the estimated structure parameters for the HN GARCH(1,1) process. Additionally, the value of the log likelihood function `$mllh` and the list `$moments` with the theoretical moments of the (unconditional) log return distribution for the estimated HN GARCH(1,1) process are returned. As input variable, `model` contains the values for the structure parameters with which the estimation procedure is started. The vector `x` contains the log return time series from which the HN GARCH(1,1) structure parameters will be estimated. If the logical flag `doprint` is set to the value `true`, the iteration path of the estimation process is printed. If the logical input variable `symmetric` is set to the value `true`, the parameter γ is set zero during the whole estimation procedure and will not be estimated.

Example: HN GARCH(1,1) Parameter Estimation - `xmpHNGmle`

Estimate the HN GARCH(1,1) structure parameters from the log return of the SPI Futures from the Sydney Futures Exchange, SFE. As a guess for the unconditional variance σ^2 , the variance of the empirical log returns of the share index is taken. The start values are $\lambda = -0.5$, $\omega = 0.1\sigma^2$, $\alpha = 0.5\sigma^2$, $\beta = 0.4$ and $\gamma = 100$. The Splus function calls take the following form

```
data <- scan(file="aofres.csv")
sigma2 <- var(data)
model <- list(lambda=-0.5, alfa=0.5*sigma2, beta=0.4, omega=0.1*sigma2, gamma=100)
hngarch.mle(model=model, data=data)
```

The results of the maximum log likelihood parameter estimation are displayed in the following table.

	symm. HN	asymm. HN	NYSE log returns
λ	3.90	2.51	-
ω	$1.15 \cdot 10^{-5}$	$2.56 \cdot 10^{-6}$	-
α	$2.59 \cdot 10^{-5}$	$1.28 \cdot 10^{-6}$	-
β	0.757	0.858	-
γ	0	58.6	-
log likelihood	13635	13676	-
persistence	0.757	0.902	-
σ^2	0.000154	0.000157	-
mean	0.00060	0.00032	0.000406
variance	0.000154	0.000157	0.000191
skewness	0.019	0.014	-6.44
kurtosis	3.398	3.37	192

■ Table: Parameters estimated from the NYSE Index data via maximum log likelihood method for symmetric and asymmetric Heston Nandi GARCH(1,1) processes. For the estimated models the theoretical values for the persistence, the unconditional variance σ^2 , the mean, the variance, the skewness and the kurtosis of the (unconditional) log return distribution are displayed. In addition the first four empirical moments are displayed.

Simulation of Asset Prices

The function `hngarch.sim(model, r=0, S0=1, innov, start.innov, doplot=T, doprint=T)` allows to simulate sample paths of a HN GARCH(1,1) stochastic process for a given set of (estimated) structure parameters. `hngarch.sim()` returns a list, containing the vector of the simulated log return time series `$x` and the vector with the associated conditional variances `$h`.

The list `model=list(lambda, omega, alpha, beta, gamma)` specifies the structure parameters of the HN GARCH(1,1) process. The interest rate term `r` is simply added to the log return time series and is set zero by default. The innovations for the artificial sample path are provided by the vector `innov`. The vector `start.innov` contains the starting innovations, which will be omitted in the output time series, but which are used, to bring the time series into a state of equilibrium. The artificial log-return time series and the associated conditional variance time series have the same length as the vector of the innovations. The initial asset price $S_{t=0}$ is specified by the variable `S0`. If the logical flag `doplot` is set true, the simulated asset prices and the percentual log returns are plotted as a function of time. Finally, by setting the logical flag `doprint` on the value true, additional information about the simulated asset price path can be provided.

Example: Simulation of HN GARCH(1,1) Price Paths - `xmpHNGsim`

Simulate HN GARCH(1,1) asset price paths and the associated log returns.


The absolute log returns of financial market time series show a characteristic long range dependence in the autocorrelation function. In contrast, the autocorrelations of the artificial absolute log returns of a HN GARCH(1,1) process decay at exponential rate and show significant values only for small lags.

It is obvious, that a stationary HN GARCH(1,1) does not fit the empirical data very well. A possible explanation for this might be, that the assumption of a stationary HN GARCH(1,1) process, describing the empirical data is less realistic, since the rules and the behavior of a financial market clearly change as time passes by. Therefore, we assume in the following, that the data are not described by a stationary HN GARCH(1,1) process during the whole time period and we allow the structure parameters to change over time: The empirical time series is divided into subintervals of equal length, and for each interval the HN GARCH(1,1) structure parameters are estimated separately. The result is a non stationary log return time series, consisting of stationary HN GARCH(1,1) processes with different structure parameters in each time interval. The autocorrelation function of the artificial log returns of such a non-stationary process shows significant values even for large lags. This illustrates, that a change of the structure parameters may be of great importance for the modelling of financial market time series, and may lead to a much better fit to the empirical data.

The figure shows the plots of the SFE SPI Futures data and its log returns in comparison to artificial log return time series which are realizations of the estimated stationary and the estimated non-stationary HN GARCH(1,1) processes.


Diagnostics - Extraction of Conditional Variances and Innovations

The Splus function `hngarch.diag(model, r=0, ht1=-1, x, doplot=T)` allows to recalculate the conditional variances and the associated innovations from a log return time series, when the structure parameters of the associated HN-GARCH(1,1) process are known. `hngarch.diag()` returns a list containing the vector with the conditional variance time series `$h` and the vector with the associated innovations `$z`, which are recalculated from the log return time series that is provided by the vector `x`. The structure parameters of the HN GARCH(1,1) process are specified in the list `model=list(alfa, beta, omega, gamma, lambda)`. Additionally an interest rate term `r` may be specified, which is simply subtracted from the log return time series `x`. In order to start the iteration to extract the conditional variances, a start value for the conditional variance `ht1` is needed. If `ht1` assumes a negative value (which is the default), the iteration starts with the value of the unconditional variance $\sigma^2 = \frac{\omega + \alpha}{1 - \beta - \alpha\gamma^2}$ used for `ht1`. The vectors of the conditional variance time series `$h` and the innovations `$z` have the same length as the




■ Figure 4.4.1: The SFE SPI Index Futures, the associated log returns and the artificial log return time series which are realizations of a simulated stationary and a simulated non-stationary HN GARCH(1,1) processes. These graphs are the output of the Splus example `xmpHNGsim.scc`.


Empirical Data



Simulated Stationary Data



Simulated Non-Stationary Data



■ Figure 4.4.2: Autocorrelations of the absolute log returns of the SFE SPI Index Futures, and of realizations of the simulated stationary and non-stationary HN GARCH(1,1) processes. The graphs are the output of the Splus example `xmpHNGsim.scc`.

vector of the log returns \mathbf{x} . If the logical flag `doplot` is set to true the time series of the log returns, the recalculated innovations and conditional variances are plotted.

The extraction of the conditional variances and the associated innovations of the log return time series can be used to test if the innovations are identical independent normally distributed random variables, which were the case if the specified model fitted the financial market data perfectly. The dependence between the innovations can be analyzed by an investigation of the autocorrelations between the innovations and with help of statistical tests.

Example: Conditional Variances and Innovations - xmpHNGdiag

Extract and plot the conditional variance time series and the associated innovations the SFE SPI Index Futures time series. Use the structure parameters of the HN GARCH(1,1) process as estimated via the maximum likelihood method. Furthermore, plot the autocorrelations of the innovations, showing that the innovations are not perfectly independent random variables. This implies that a HN GARCH(1,1) process with normally distributed innovations is not a perfect model for the SFE SPI Index Futures data.

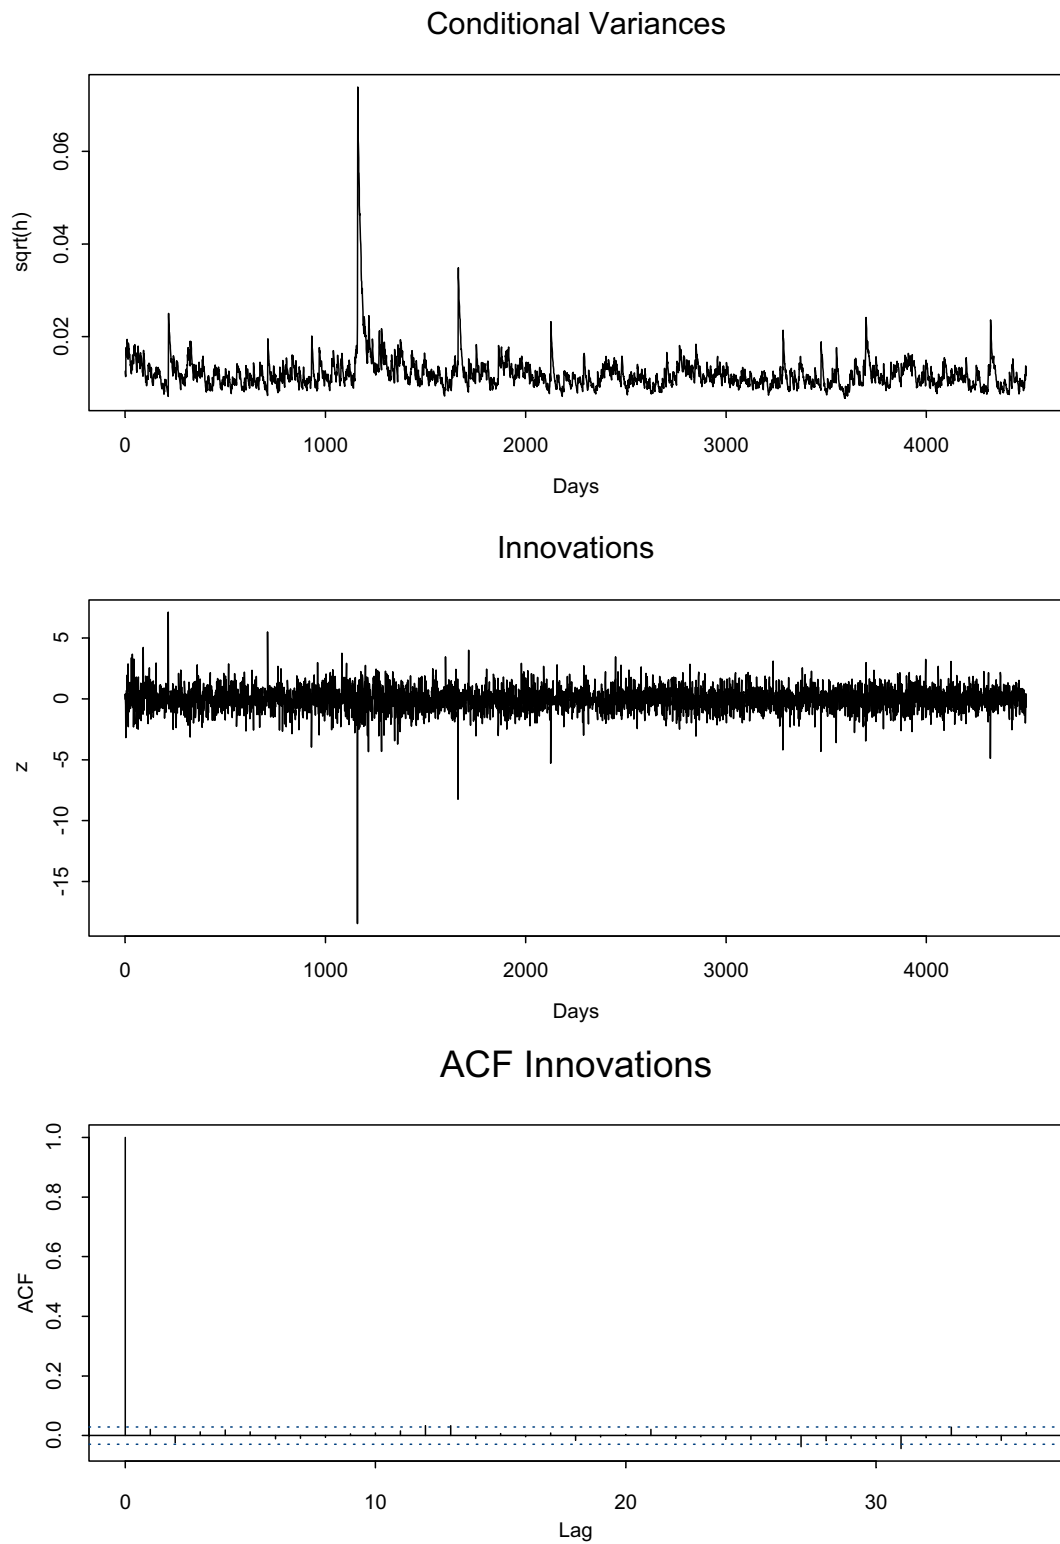
Moments of the Log Return Distribution

The Splus function `hngarch.mom(model, r=0)` calculates the first four moments of the unconditional log return distribution for a stationary HN GARCH(1,1) process with standard normally distributed innovations. `hngarch.mom()` returns a list with the theoretical values for the `$mean`, the `$variance`, the `$skewness` and the `$kurtosis` of the (unconditional) log return distribution. We have also access to the `$persistence` of the corresponding HN GARCH(1,1) process and to the values for $E[\sigma_t^2]$ (`$meansigma2`), $E[\sigma_t^4]$ (`$meansigma4`), $E[\sigma_t^6]$ (`$meansigma6`) and $E[\sigma_t^8]$ (`$meansigma8`), which are needed for the computation of the moments of the unconditional log return distribution. The only arguments are the risk free interest rate \mathbf{r} and the structure parameters of the HN GARCH(1,1) process, which are specified in the model list `model=list(alfa, beta, omega, gamma, lambda)`.

The Splus function `hngarch.momdiff(model, r=0)` returns a list with the associated partial derivatives of the persistence, the unconditional variance (and the other moments of the conditional variance), the mean, the variance, the skewness and the kurtosis of the log returns, with respect to the structure parameters λ , α , β , ω and γ of the HN GARCH(1,1) process. The HN GARCH(1,1) structure parameters for which the partial derivatives shall be computed are specified in the list `model`. The partial derivative of the kurtosis with respect to the parameter α , for example can be accessed with `$dkurtosisalfa` `hngarch.momdiff()` can be used in order to compute the derivatives of the HN option price with respect to the physical parameters, such as the unconditional variance, the persistence or the kurtosis.

Pricing of European Call Options

The Splus function `HNGOption(CallPutFlag, model, S, X, Time, r=0, ht1)` evaluates the prices of call and put options according to the Heston-Nandi GARCH(1,1) model. `HNGOption()` returns the price for the selected option styke. The arguments are the `CallPutFlag` taking the string values "c" or "p", the HN GARCH(1,1) `model` list the asset price \mathbf{S} , the strike price \mathbf{X} , the `Time` to maturity, the risk free interest rate \mathbf{r} , and the volatility `ht1`. The HN prices for put



■ Figure 4.4.3: The conditional variance time series and the associated innovations recalculated from the NYSE Index log return time series. The model parameters were estimated with the max log likelihood method in Example 1. The figure on the bottom shows the autocorrelations of the innovations time series.

options are evaluated by this function using the put-call parity ³. The corresponding functions are

Example: Scaling Properties of Option Prices - xmpHNGscaling

This Splus script demonstrates the usage of the option pricing functions by showing the scaling properties of put option prices. An analysis of the Black-Scholes option pricing formula reveals that the prices of Black-Scholes options depend only on the arguments $\sigma^2(T-t)$ and the ratio $S/Xe^{-r(T-t)}$. By doubling the volatility, and dividing the time to maturity by the factor 4, we do not change the price of a BS option. The Figure shows that a similar result holds for the HN option prices. It has to be noted, that the doubling of the volatility not only corresponds to the doubling of the conditional variance `ht1` at time t but also corresponds to a doubling of the unconditional variance σ^2 , which plays the same role in the HN GARCH model as the squared volatility in the Black Scholes model. Therefore, we chose the structure parameters for the HN GARCH(1,1) model as fixed functions of the unconditional variance ($\alpha = 0.5\sigma^2$, $\beta = 0.4$, $\gamma = 0.1/\sigma$, $\omega = 0.595\sigma^2$, $\lambda = -0.5$) and scale them only by scaling the value for the unconditional variance.

The Greeks of an European Call Option

The Greeks denote the partial derivatives of the option price with respect to the asset price S (the delta), the volatility σ (the vega), the time t (the theta) and the interest rate r (the rho). The gamma denotes the option price derived two times with respect to the asset price S . The Splus function `HNGGreeks(CallPutFlag, Selection, model, S, X, Time, r=0, ht1)` returns the value of a HN GARCH(1,1) option Greek of a call or put option, which is selected in the argument `Selection`. The other arguments are the same as for the Splus function `HNGOption()`.

Example: HN Greeks of a Call Option - xmpHNGGreeks

The Splus script `Example6.scc` plots the Greeks of a Heston Nandi call option for different asset/strike price ratios and for different times to maturity.

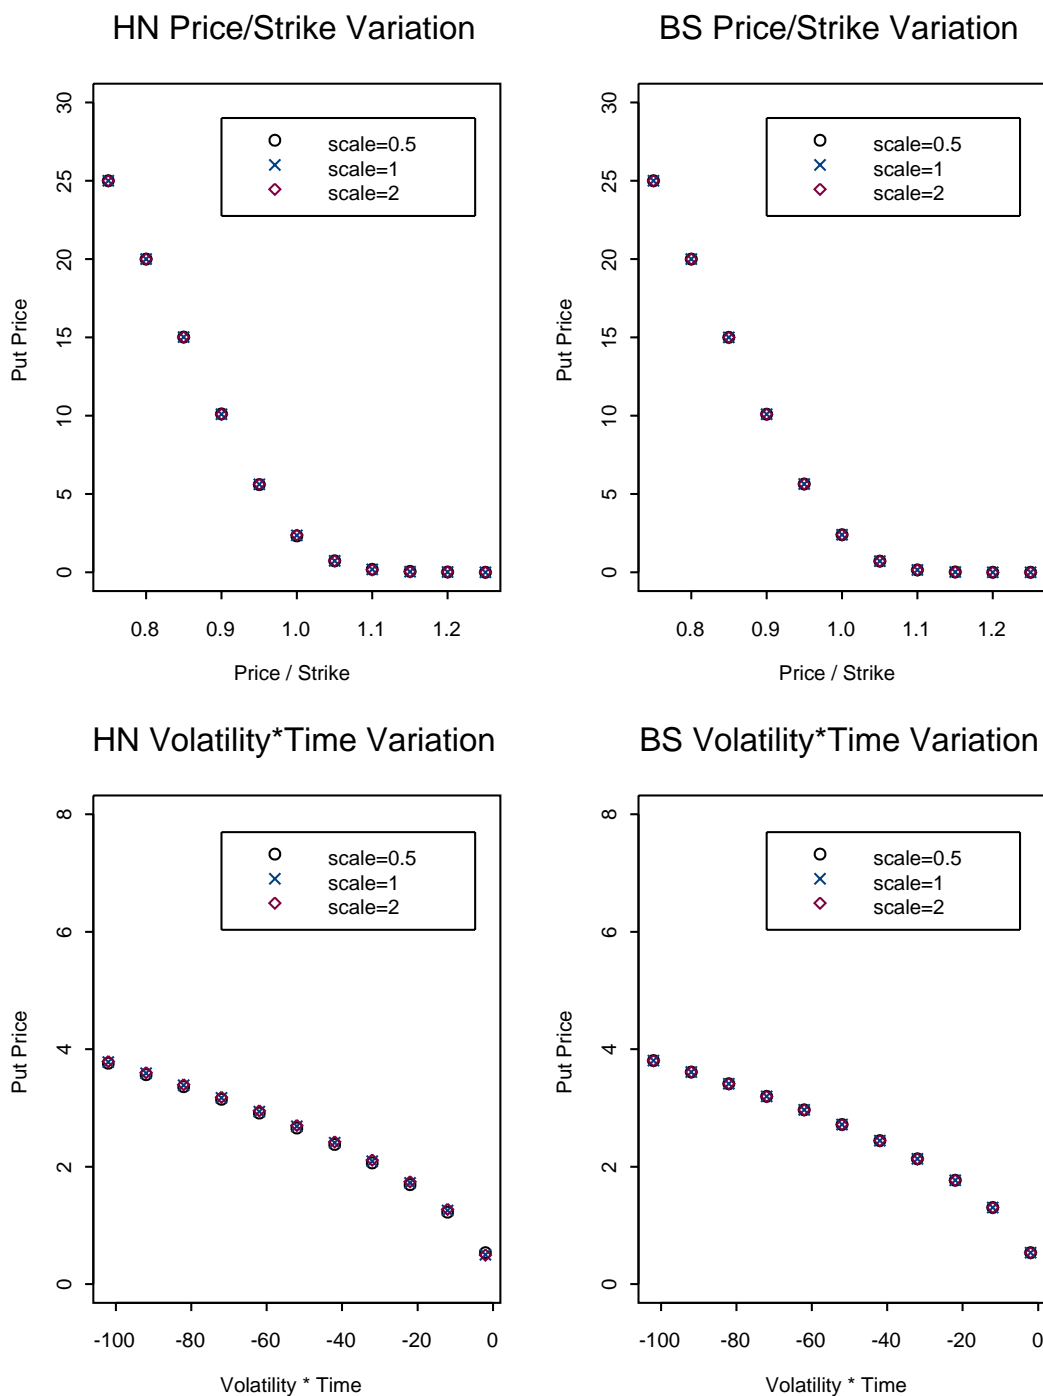
Black-Scholes Implied Volatility Derived from Market Prices

The Splus function `bs.imv(CallPutFlag, price, S, X, Time, r=0, ...)` allows to extract the Black-Scholes implied volatility from the market price c of an European call option, the asset price S , the strike price X and the time remaining to expiration ($T-t$). The implied volatility is computed by solving for the zero in the difference between the BS price and the market price of the option, with the volatility as the variable. `bs.imv()` returns the daily squared BS implied volatility of a call option with the market price c , the risk free interest rate r , the asset price S , the strike price X and the time to maturity `time`.

Example: The Smile Effect

In the Splus script `Example8.scc` we first compute the HN call option prices for four different sets of HN GARCH(1,1) structure parameters. Then the annualized Black-Scholes implied volatility is computed with the help of the `bs.imv` function. The BS implied volatility of the HN GARCH(1,1) option prices as a function of S/X , shows the same "smile pattern" as this is often observed for real option prices.

³Put-Call Parity: $c + Xe^{-r(T-t)} = p + S$ where c is the price of an European call option and p is the price of an European put option. X denotes the strike price and S denotes the asset price



■ Figure 4.4.4: The BS and HN put option prices as functions of S/X and the product $\sigma^2(T-t)$. In the HN GARCH model σ^2 denotes the unconditional variance. In each plot three different scalings are displayed. The time to maturity is multiplied by the factor $scale$ and the squared volatility is divided by the factor $scale$, such that the product $\sigma^2(T-t)$ is a constant. The put option in the lower row is at the money.

Root Mean Square Parameter Estimation and Option Misspricing

We could price an option by estimating the HN GARCH structure parameters from historical asset price data via the maximum log likelihood method, and then using these parameter estimates to evaluate the HN GARCH option price. Similarly we could estimate the BS volatility from historical asset prices and use this estimate to evaluate the BS option price. Since the price of an option reflects the expectations about the future evolution of the asset prices, the pricing of an option that is based solely on historical information could be not suitable. Therefore it is more appropriate to estimate the needed quantities from the historical option prices, which to some extent contain information about the expectations about the future behavior of the asset prices.

Since the BS and the HN GARCH model have both closed form solutions for the option prices, it is quite natural to estimate the underlying structure parameters by minimizing the difference between the market prices and the model prices. The difference may be expressed as sums of volume weighted squared errors (L2 norm). There have been implemented two Splus functions for the root mean square error (RMSE) estimation, which allow for the BS model to find the optimal value for the volatility $\Theta = \sigma$, and for the HN GARCH(1,1) model to find the optimal structure parameters $\Theta = (\alpha, \beta, \omega, \gamma, \lambda)$. The estimations are done by minimizing the root mean squared error

$$RMSE = \frac{\sum_{Options} (c_{Market} - c_{Model})^2 Volume}{\sum_{Options} Volume}. \quad (4.70)$$

Options is a set of traded call options with the prices c_{Market} , the asset prices S , the strike prices X and the times to maturity $(T-t)$. c_{Model} are the associated call option prices according to the BS or the HN GARCH model. The (relative) misspricing is defined as the difference between c_{Market} and c_{Model} divided by the market call option price.

```
bs.rms (c, r=0, S, X, time, volume, tolerance=1e-8)
```

returns a list containing the estimated volatility `$sigma2` of the BS model, the RMS error `$rmse` and the vector of the misspricings `$misspricings`. The arguments of this function are the vector `c` with the market prices of the selected call options, the risk free interest rate `r`, the vector `S` with the asset prices, the vector `X` with the strike prices, the vector `time` with the times to maturity (in days) and the vector `volume` with the traded volume of the options. The tolerance to which accuracy the RMS error will be determined is specified in the argument `tolerance`.

For the HN GARCH(1,1) model additional arguments are needed: the starting values for the structure parameters and the time series with the prices of the underlying. The latter argument is needed to extract the conditional variances for the estimated model. The conditional variances need not to be estimated, since they can be recalculated from the asset price history.

```
hngarch.rms (model, c, r=0, X, time, volume, U, utime, tolerance=1e-8, doprint=T)
```

returns a list with the estimated HN GARCH(1,1) structure parameters `$parameters = list(lambda, omega, alfa, beta, gamma)`, the root mean square error `$rmse`, the vector `$misspricings`

with the misspricings, and the vector `$ht1` with the conditional variances. The needed arguments include the list `model = (lambda, omega, alfa, beta, gamma)`, which specifies the starting parameters for the RMS error estimation, the vector `c` with the market prices of the selected call options, the risk free interest rate `r`, the vector `S` with the asset prices, the vector `X` with the strike prices, the vector `time` with the times to maturity (in days), and the vector `volume` with the traded volume of each option. `U` is the vector with the prices of the underlying and `utime` is the vector with the associated days until expiration. The tolerance to which accuracy the RMS error will be determined is specified in the argument `tolerance`.

Example: HN GARCH(1,1) Structure Parameters from SFE Future Calls - `xmpHNGrms`

This Splus script applies the functions `bs.rms` and `hngarch.rms` on traded call options with the SFE future contracts as the underlying asset. The prices of the SFE futures, the traded call options and the relative misspricing for the estimated BS and HN GARCH(1,1) models are shown in Figure 4.4.2. The most traded call options, i.e the options with the largest trading volume, are at the money. Therefore the misspricing is small for the at the money options, since then the root mean square error is small. The misspricing for the traded call options with the largest trading volume is smaller for the HN GARCH model than for the BS model, because there are more degrees of freedom in the HN GARCH(1,1) model, such that is able to fit the empirical data better than the BS model. For the options that are not at the money, however, the misspricing is smaller for the BS model as it can be verified in Figure 4.4.2. Note furthermore, that the estimation of the HN GARCH(1,1) structure parameters directly from the SFE future prices via the maximum log likelihood method leads to results which differ a lot from the root mean squared error parameter estimation indirectly from the option prices.

Estimation method	λ	ω	α	β	γ
root mean square error	0.49	$3.3 \cdot 10^{-5}$	$1.9 \cdot 10^{-9}$	0.46	3005
max log likelihood	7.07	$4.6 \cdot 10^{-6}$	$3.3 \cdot 10^{-7}$	0.01	1660

■ HN GARCH(1,1) structure parameters estimated directly from the SFE future contract prices via maximum likelihood method and estimated indirectly from the associated European call option prices via RMS error minimization.

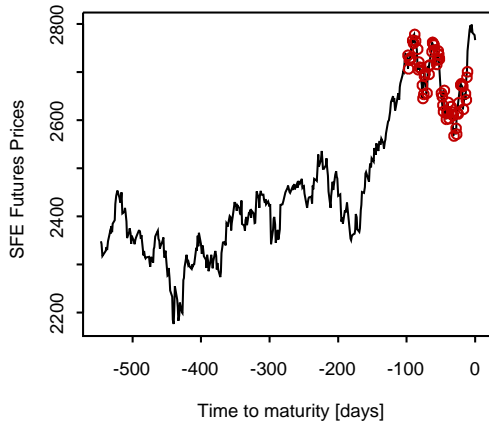
Hedging of European Call Options

In the Black-Scholes (implied volatility) delta hedging strategy, the hedge ratio h is set equal to the BS delta of the derivative. The volatility σ , which is needed for the computation of the BS delta, is estimated from the current call option prices, when the hedge ratio is updated.


```
bs.hedge(S, c, time, X, r=0, sigma2=-1
```

returns a list containing the vector `$hedgeratio` with the BS (implied volatility) delta hedge ratios at the times when the hedge ratio is updated. The times when the hedge ratio is updated are specified in the input vector `time`. Additional arguments are the vector `S` and the vector `c` which contain the asset and the call option prices at the times when the hedge ratio is updated. The argument `X` denotes the strike price of the hedged call option and the argument `r` denotes the risk free interest rate. If the argument `sigma2` assumes a negative value (which is the default), the implied volatilities from the option prices in `c` are used for the computation of the hedge ratios. Otherwise the squared volatility is assumed as a constant with the value `sigma2`. Note that the estimation of the BS implied volatility becomes numerically unstable for times near to maturity, such that a constant volatility (which has to be estimated somehow) might be useful.


Futures and traded SFE Call Options




Traded SFE Call Options




BS Call Option Misspricing



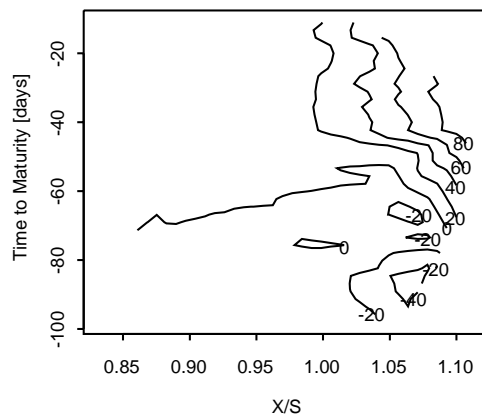
BS Call Option Misspricing



HN Call Option Misspricing



HN Call Option Misspricing



■ The prices of the SFE future contracts, the traded call options and the percentual misspricings between the traded and the BS/HN call options.

Besides the hedge ratios the output list contains the hedge errors and the gains from trade during the time interval between two successive dates when the hedge ratio is updated. The associated vectors can be accessed by `$hedgeerror` and `$gainsfromtrade`.

Example: Hedge Performance of Discrete Time Black-Scholes Delta Hedging - xmpBSHedge

This Splus script simulates 1000 asset price paths of length 100 days, which are realizations of a Brownian motion process with a volatility of 16% per trading year. For each path the function `bs.hedge` is used to determine the BS (implied volatility) delta hedge errors during one day, which is the time interval between the updating of the hedge ratio. The variance of the hedge errors and the average absolute hedge errors are then evaluated for each of the 100 days before maturity. When the hedge ratio is updated less frequently, the hedge performance of the BS delta hedging strategy gets worse.

The Heston-Nandi hedging strategy uses the hedge ratio (4.69), which differs from the HN delta of the call option especially for γ parameters different from zero. The conditional variances for the computation of the hedge ratio are recalculated according to the relation

$$\sigma_{t+1}^2 = \omega + \beta\sigma_t^2 + \alpha \frac{(\log S_t - \log S_{t-1} - r - \lambda\sigma_t^2 - \gamma\sigma_t^2)^2}{\sigma_t^2}. \quad (4.71)$$

Unlike as in the BS delta hedging strategy the volatility can be recalculated directly from the asset prices. Instead the structure parameters of the HN GARCH(1,1) model have to be provided by an estimation.

```
hngarch.hedge(S, c, X, U, time, ht1=-1, r=0)
```

returns a similar output list for the HN hedging strategy as the function `bs.hedge` for the BS delta hedging strategy. The structure parameters for the HN GARCH(1,1) model are specified in the list `model`. For the recalculation of the conditional variances, the asset price time series `U` with the prices at each discrete time point from the day on when the hedging is started has to be provided. The value for the conditional variance at the day when the hedging is started is provided by the argument `ht1`. If `ht1` assumes a negative value, the start value for the conditional variance is set equal to the unconditional variance of the specified HN GARCH(1,1) process.

Example: Hedge Performance of the Heston-Nandi Hedging Strategy - xmpHNGhedge

The Splus script `Example11.scc` demonstrates the application of the function `hn.hedge`, by showing that the HN hedging strategy performs better than the BS (implied volatility) delta hedging strategy for options traded in an artificial financial market which behaves according to the HN GARCH(1,1) model.

Notes and Comments

In section 1, most of the Splus functions were derived from the Excel Spreadsheets as presented in the book of Haug. For the software implementation, done by D. Würtz, in most cases only minor changes were required to make them ready for Splus.

The Exotic options were also implemented from the book of E.G. Haug (1997).

The software package for pricing Heston Nandi Options was written by D. Würtz, and some functions were added by R. Angliker. These Splus functions rely on several Fortran77 routines: **DUMINF** - BFGS optimization routine to perform the max log likelihood estimation of the HN GARCH(1,1) structure parameters; **DQAGI** - Semi-infinite integration routine to evaluate the integrals in the HN option pricing formulas; **DFZERO** - bisection/secant algorithm in order to search for the zero to solve for the BS implied volatility; **DNLS1E** - Levenberg-Marquard routine to perform the root mean square error estimation of the HN GARCH(1,1) structure parameters; **CN** to calculate the cumulative normal distribution function in the pricing formula for the BS options; **CDEXP2**, **CDLOG2** substitutions for the double precision non-standard Fortran77 routines for portability reasons.

4.5 Monte Carlo Simulations of Options

4.5.1 The Monte Carlo Approach

A Monte Carlo Simulation for pricing options approximates the expected derivative's cash flows with a simple arithmetic average of the cash flows over a finite number of simulated price paths. Consider, for example, the simple case of a Plain Vanilla European call option. The price of this option is expressed as the expectation value of its payoff, discounted at the risk-free interest rate.

$$c = e^{-r(T-t)} E[\max(S_T - X, 0)] .$$

If we simulate enough possible stock prices, say $N = 100'000$, and take the average of the returns of the option, we expect to make a good approximation for the real price.

$$c \simeq e^{-r(T-t)} \sum_{i=1}^N \max(0, S_T^{(i)} - X) / N , \quad (4.72)$$

where

$$S_T^{(i)} = S \exp \left[\left(r - \frac{1}{2} \sigma^2 \right) T + \sigma T^{\frac{1}{2}} \varepsilon^{(i)} \right] \quad \text{and} \quad \varepsilon^{(i)} \sim \phi[0, 1] , \quad (4.73)$$

and the $i = 1, 2, \dots, N$ enumerates the configurations.

The Monte Carlo procedure is, in fact, not adapted to price path-independent options, where the return value of the option is only determined by the final value of the price. Compared to other procedures, Monte Carlo simulation is numerically efficient when there are several variables to define the return value of a price path, that is a path-dependent option. The option price is then the expectation value of the return of these paths, that is a multi-dimensional integral over the variables that characterize a path. Consider the multi-dimensional integral

$$I = \int_0^1 d^d x f(x) . \quad (4.74)$$

In general, traditional approaches to multi-dimensional integration grow exponentially with the dimension of the space, d , in which the integral is being performed. So, to obtain a certain accuracy, the required CPU time T_{CPU} to run the algorithm, scales as

$$T_{CPU} \sim \epsilon^{d/\alpha} , \quad (4.75)$$

where ϵ is a constant, and α is determined by the particular integration “rule” being used (*i.e.* first-order, second-order etc.). In contrast, for stochastic techniques, like Monte Carlo, the error is due to random sampling over the integration region, expecting for the required CPU time

$$T_{CPU}^{MC} \sim \frac{1}{N} \sum_{i=1}^N f(x^{(i)}) . \quad (4.76)$$

The resulting error scales much more favorably, as $1/\sqrt{N}$. Actually the required computation time doesn’t explicitly depend on the dimension of the integral. The variance is as important as the mean in stochastic sampling; without it, one cannot assign a significance to any result: the way the estimate $\langle f \rangle$ approaches the true value of the integral is determined by the variance. The fact that the variance is independent of the number of dimension is the foundation of Monte Carlo integration and can be shown formally by the central limit theorem, which states that the distribution of averages approaches a normal distribution, whose spread around the central value, determined by the square root of the variance, is proportional to $1/\sqrt{N}$. We just recall here the formulation of the central limit theorem:

Central Limit Theorem

If X_1, X_2, \dots are a set of random independent variables identically distributed with mean μ and variance σ^2 , then the distribution of


$$\frac{X_1 + X_2 + \dots + X_N - N\mu}{\sigma\sqrt{N}} \quad (4.77)$$

tends to a standard normal distribution when $N \rightarrow \infty$.

This result can expressed as:

$$\lim_{N \rightarrow \infty} P \left\{ \frac{X_1 + \dots + X_N}{N} - \mu \leq a \right\} = \frac{1}{\sqrt{2\pi}(\sigma/\sqrt{N})} \int_{-\infty}^a e^{-\frac{N}{2\sigma^2}x^2} dx, \quad (4.78)$$

showing that the expression $\frac{X_1 + \dots + X_N}{N}$ tends to be normally distributed with mean μ and variance σ^2/N . This can be, of course, applied to equation (4.76) to conclude that a Monte Carlo estimator converges to the true value with variance proportional to $1/N$ thus a spread around the mean (error) proportional to $1/\sqrt{N}$ with a factor of proportionality equal to σ , the square root of the variance of the distribution $f(x)$. Figure 4.5.1 shows, for a Plain Vanilla European option, the relative error, as function of the Monte Carlo steps with a squared root scale for the x-axis. The relative error is defined here as e/c where c is the Black-Scholes price and $e = 2\sigma$. This guarantees that more than 95% of the estimation are include in $c \pm e$.



■ Figure 4.5.1: Monte-Carlo simulation: Estimation of the relative error as function of N on the basis of 500 independent price estimations of a Plain Vanilla option with $S = 100$, $X = 100$, $\sigma = 0.5$, $t = 1$, and $r = 0$.

4.5.2 Monte Carlo Estimators of the Greeks

Over the last decade, a variety of direct methods have been developed for estimating partial derivatives by simulation. Direct methods compute a derivative estimate from a single simulation, and thus do not require re-simulation at a perturbed parameter value. Under appropriate conditions, they result in unbiased estimates of the derivatives themselves, rather than of a finite-difference ratio. Our discussion focuses on the use of *pathwise* derivatives as direct estimates, based on a technique generally called *infinitesimal perturbation analysis*, see Glassermann (1991).

The pathwise estimate of the true Black-Scholes Delta dc/dS is the derivative of the sample price $c^{(i)}$ with respect to S_0 . More precisely, it is

$$\frac{dc^{(i)}}{dS} = \lim_{\delta \rightarrow 0} \frac{c^{(i)}(S + \delta) - c^{(i)}(S)}{\delta},$$

provided the limit exists with probability 1. If $c^{(i)}(S)$ and $c^{(i)}(S + \delta)$ are computed from the same random number $\varepsilon^{(i)}$, then, provided $S_T^{(i)} \neq X$, we have

$$\frac{dc^{(i)}}{dS} = \frac{dc^{(i)}}{dS_T^{(i)}} \frac{dS_T^{(i)}}{dS} = e^{-r} \mathbf{1}_{\{S_T > X\}} \frac{S_T^{(i)}}{S}, \quad (4.79)$$

Here we used equation (4.73) to get

$$\frac{dS_T^{(i)}}{dS} = e^{(r - \frac{1}{2}\sigma^2) + \varepsilon^{(i)}} = \frac{S_T^{(i)}}{S}$$

and

$$\frac{dc^{(i)}}{dS_T^{(i)}} = e^{-r} \frac{d}{dS_T} \max(0, S_T - X) = \begin{cases} e^{-r} & , S_T > X, \\ 0 & , S_T < X. \end{cases}$$

As S_T , Cc fails to be differentiable; however, since this occurs with probability zero, the random variable $\frac{dc^{(i)}}{dS}$ is almost surely well defined.

The pathwise derivative $\frac{dc^{(i)}}{dS}$ can be thought as limiting case of the common random numbers finite-difference estimator, in which we evaluate a limit analytically rather than numerically. It is a direct estimator of the option delta because it can be computed directly from a simulation, starting at S without the need for a separate simulation at a perturbed value S . The question remains whether this estimator is unbiased; that is whether

$$E \left[\frac{dc^{(i)}}{dS} \right] = \frac{dc}{dS} \equiv \frac{d}{dS} E \left[c^{(i)} \right] . \quad (4.80)$$

The bias of the pathwise estimate thus reduces to the interchangeability of derivative and expectation. The interchange is easily justified in this case; see Broadie and Glassermann (1993) for this example and conditions for more general cases. Applying the same reasoning used above, we obtain the following pathwise estimators of Vega for the Black-Scholes price:

$$Vega(dc/d\sigma) : e^{-r} \mathbf{1}_{S_T > X} \frac{S_T}{\sigma} \left(\ln(S_T/S) - (r - \frac{1}{2}\sigma^2)T \right) \quad (4.81)$$

Each of these estimators is unbiased. Of course, Monte Carlo estimators are not required for these derivatives because closed-form solutions are available for each. The Black-Scholes setting is useful for illustration, but the utility of the technique rests in the applicability to more general models. Broadie and Glassermann (1993) have derived and studied pathwise estimates, both theoretically and numerically, for Asian options and a model with stochastic volatility. For example, the Asian option delta estimate is simply

$$e^{-r} \frac{\bar{S}}{S} \mathbf{1}_{\{\bar{S} > X\}} , \quad (4.82)$$

where \bar{S} is the average asset price used to determine the option payoff. Evaluating this expression takes *negligible* time compared with a re-simulation to estimate the option price from perturbed initial stock prices. The pathwise estimate is thus both more accurate and faster to compute than the finite-difference approximation. These advantages extend to a wide class of problems.

4.5.3 Variance-Reduction Techniques

Computer time efficiency is the crucial point for fast and reliable Monte Carlo simulations. Thus we present in the following techniques to reduce the variance in Monte Carlo Simulations. The advantage of variance reduction, which means obtaining faster convergence with less sample points, is obviously desirable for Monte Carlo simulations, but sometimes overlooked.

Random Number Generator Efficiency

We must, however, not forget to take into account the computation time required to generate a random sequence. If a certain way of generating random numbers allows to obtain variance reduction, it is only efficient when the computer time necessary to generate these random numbers is sufficiently low. More precisely, suppose that we have a choice between two types of random generators for a Monte Carlo simulation. Monte Carlo samples of N points approximate the real value to estimate with variances σ_1^2 and σ_2^2 , respectively. If we denote with b_1 the time to generate a random number in the first sequence and b_2 the time to generate a random number in the second, the first method is better than the second only when

$$\sigma_1^2 b_1 < \sigma_2^2 b_2 . \quad (4.83)$$

The value $e \equiv \sigma^2 b$ is sometimes referred as the *efficiency* of the procedure, their ratios e_1/e_2 , will allow us to make comparisons.

Antithetic Variates

One of the simplest and most widely used techniques in financial pricing problems is the method of *antithetic variates* as variance-reduction technique. Consider the problem of computing the Black-Scholes price of an European call option. Independent replications of the terminal stock prices are generated from formula (4.73). In this context, the method of antithetic variates is based on the observation that if ε is a standard normal distribution, then so does $-\varepsilon$. Similarly

$$c^{(-)} = e^{-r(T-t)} \sum_{i=1}^N \max \left(0, S_T(-\varepsilon^{(i)}) - X \right) / N \simeq c,$$

is an unbiased estimator of the option price, as is therefore

$$c_{AV} = \frac{1}{2}(c^{(+)} + c^{(-)}) . \quad (4.84)$$

We can now compare the efficiencies as follows. Because $c^{(+)}$ and $c^{(-)}$ have the same variance,

$$\text{var} \left[\frac{1}{2} (c^{(+)} + c^{(-)}) \right] = \frac{1}{2} \left\{ \text{var} [c^{(+)}] + \text{cov} [c^{(+)}, c^{(-)}] \right\} . \quad (4.85)$$

Thus we have

$$\text{var} [c_{AV}] \leq \text{var} [c^{(+)}] \quad \text{if} \quad \text{cov} [c^{(+)}, c^{(-)}] \leq \text{var} [c^{(+)}] .$$

However, c_{AV} uses twice as many replications as $c^{(+)}$, so we must account for differences in computational requirements. If generating the ε takes a negligible fraction of the work per replication, then the work to generate c_{AV} is roughly double the work to generate $c^{(+)}$. Thus for antithetics to increase efficiency, we require

$$2\text{var} [c_{AV}] \leq \text{var} [c^{(+)}] , \quad (4.86)$$

which in light of equation (4.85), simplifies to the requirement that $cov [c^{(+)}, c^{(-)}] \leq 0$. That this condition is met is easily demonstrated. The function

$$c^{(+)} = g(\varepsilon^{(1)}, \varepsilon^{(2)}, \dots, \varepsilon^{(N)}) = e^{-r(T-t)} \frac{1}{n} \sum_{i=1}^N \max \left(0, S_T(-\varepsilon^{(i)}) - X \right)$$

is the composition of the mapping from the ε 's to the stock prices and from the stock prices to the discounted option payoff. As the composition of two increasing functions, g is monotone, so we find according to Barlow and Proschan (1975)

$$E[g(\varepsilon)g(-\varepsilon)] \leq E[g(\varepsilon)]E[g(-\varepsilon)], \text{ (where } \varepsilon = (\varepsilon^{(1)}, \varepsilon^{(2)}, \dots, \varepsilon^{(N)}) \text{)} \quad (4.87)$$

i.e., $Cov[c^{(+)}, c^{(-)}] \equiv E[g(\varepsilon)g(-\varepsilon)] - E[g(\varepsilon)]E[g(-\varepsilon)] \leq 0$, and we may conclude that antithetics improve Monte Carlo simulations.

This argument can be adapted to show that the method of antithetic variates increases efficiency in pricing a European put and other option that depend monotonically on inputs (e.g. Asian options). On the other hand, the notable departure from monotonicity in some barrier options (e.g. down-and-in call) suggests that the use of antithetics in pricing these options may sometimes be less effective.

Quasi Monte Carlo Methods: Low-Discrepancy Sequences

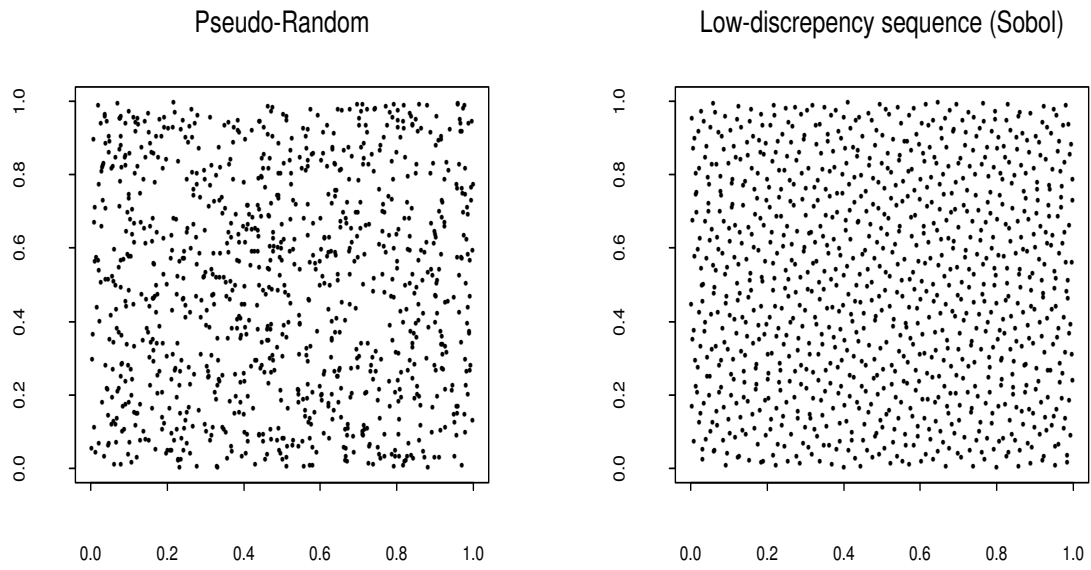
Quasi Monte Carlo Methods use deterministic sequences, called low-discrepancy sequences, which are in some sense well distributed. The main advantage we expect is that we achieve a faster convergence compared with Monte Carlo simulations based on random numbers.

To simulate a price path using low-discrepancy sequences, we first generate numbers θ in the interval $[0, 1)$, and then transform it into normally distributed numbers. For each time interval i , a new ε_i is generated, and the corresponding asset price follows from equation (4.73).

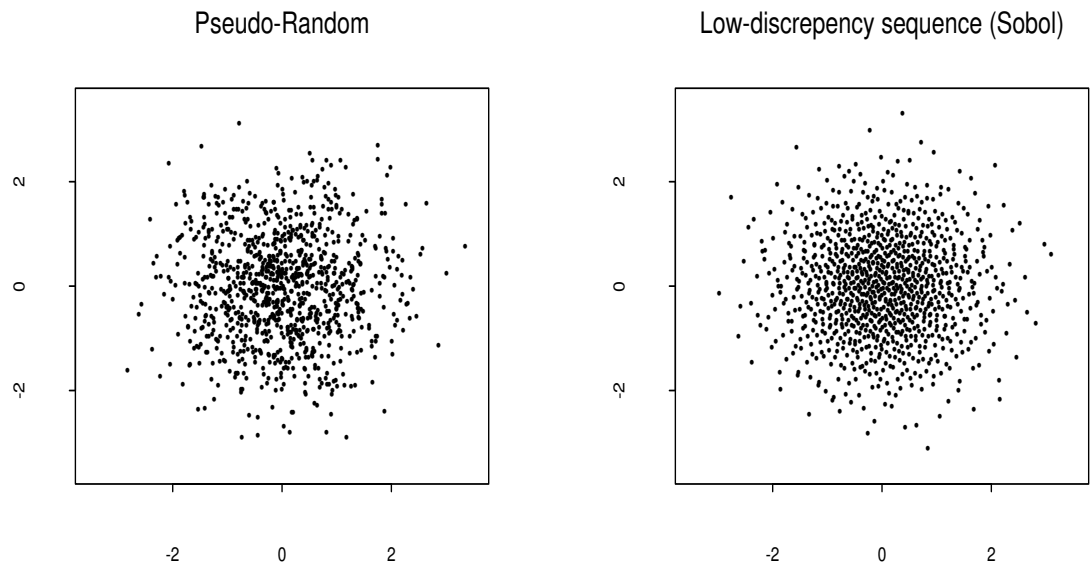
In the terminology of low-discrepancy sequences, the number of time intervals is referred to as the number of dimensions of the simulations. Each path being parameterized with n θ 's, a path is thus a point in the n dimensional "space of the θ 's" and the price, proportional to the expectation value of the option, is an integral over this space.

Very important for Monte Carlo simulations is the ability to generate "good" ε 's which means that, for each time step, the distribution of simulated ε 's across all price paths should closely approximate $\phi(0, 1)$. The fact that $\varepsilon = N_{(0,1)}^{-1}(\theta)$, reduces the problem of getting "good" ε 's to a problem of generating sequences of θ 's "well" uniformly distributed over the hypercube $[0, 1)^n$. Figure 4.5.3 simulates an uniform distribution of 1000 points over a two-dimensional space with a "standard" (pseudo) random generator. Figure 4.5.4 shows to corresponding two dimensional normal distribution. We observe that certain points are clustered together, which also implies that some regions will have few points.

In general, the more evenly the points are distributed throughout the domain, the more accurate the simulation. This is the approach of low-discrepancy sequences. The points are "deterministically" rather than randomly chosen to fill the hypercube uniformly, thereby minimizing clustering and improving accuracy.



■ Figure 4.5.3: Pseudo-Random versus Low-Discrepancy Sequences.



■ Figure 4.5.4: Normal Distributions of Pseudo-Random versus Low-Discrepancy.

Recall that a Monte Carlo simulation using random sequences generates error bounds that are of the order $1/\sqrt{N}$, where N is the number of simulated paths. Also the error bound is independent of dimensionality. Low-discrepancy sequences, on the other hand, generate an upper bound to the error that has been estimated by the Koksma-Hlawka inequality. They find an upper bound to the error of the order of $(\log N)^n/N$, where n is the number of dimensions (i.e. time intervals). This can be expressed as follow.

Kosma-Hlawaka Theorem:

Let $I^d = [0, 1]^n$ and let f have bounded variation on $[0, 1]^n$ in the Hardy-Krause⁴ sense. Then for any $x_1, x_2, \dots, x_N \in I^n$ we have

$$\left| \frac{1}{N} \sum_{k=1}^N f(x_k) - \int_{I^n} f(u) du \right| \leq C(f) \frac{(\log N)^n}{N} \tag{4.88}$$

where $C(f)$ is a constant.

The error bound provided by this theorem, while it is of theoretical interest, is unfortunately of little help in most practical situations. The constant $C(f)$ is very difficult to estimate accurately in high dimensions and “often extremely large”, see Spanier and Maize(1994).

Studies using low-discrepancy sequences in finance applications find that the errors produced are substantially lower than the corresponding errors generated by standard Monte Carlo simulations. Three low-discrepancy sequences are the Halton, Faure and Sobol sequences. We will use the Sobol sequences because their applications in finance have given the best results in the paper of Boyle, Broadie, Glasserman, (1995). It is actually the most popular low-discrepancy sequence.

Example: Sobol Random-Generator

We present here the method proposed by Antonov and Saleev’s (1979) which is an improvement of the original Sobol method (about 20 % faster). The construction of a Sobol sequence follows a three-step procedure.

Step 1: Generate a set of odd integers m_i , for $i = 1, 2, \dots, [\log_2 N] + 1$, that satisfy the condition: $0 < m_i < 2^i$, where N is the number of price paths, and $[\log_2 N]$ is the bigger integer smaller or equal than $\log_2 N$. $M = [\log_2 N] + 1$ is the maximum number of digits in the expansion of N in base 2. This step entails a quite complicated recursive procedure:

We start with a series of integers h_1, h_2, h_3, \dots where h_j is either “0” or “1”. The h_j are the coefficients of a *primitive polynomial* modulo 2

$$P(x) = \sum_{i=0}^q h_i x^i, \quad \text{where} \quad h_q = h_0 = 1, \tag{4.89}$$

and q is the dimension of the polynomial. The theory of primitive polynomial modulo 2 is a topic in number theory that we will not discuss here. It suffices here to say that there are special polynomials among those whose coefficients are zero or one. An example is $x^9 + x^4 + x^3 + x + 1$ which we can abbreviate with $(q, p) = (9, 13)$, where the bits of p in base 2 indicates the values of h_i ’s, $0 < i < q$, in the following way:

$$h_i = \left[\frac{p}{2^{i-1}} \right] \text{ mod } 2. \tag{4.90}$$

⁴For a more complete discussion of the Hardy-Krause definition variation and details on this theorem see Niederreiter (1992).

The brackets “[]” are defined as above. In our example; $p = 13 = 000001101$ indicates $h_1 = h_3 = h_4 = 1$ and $h_2 = h_5 = h_6 = h_7 = h_8 = 0$.

The first q odd integers m_1, m_2, \dots, m_q must be supplied. They can be chosen freely, provided they satisfy that m_i is odd and $0 < m_i < 2^i$. Next we generate the set of the other m_i for $q < i \leq M$ using the coefficient of the primitive polynomial and a recursive relationship for $i > q$:

$$m_i = m_{i-q} \oplus 2^q m_{i-q} \oplus 2^{q-1} h_1 m_{i-q+1} \oplus \dots \oplus 2^2 h_{q-2} m_{i-2} \oplus 2 h_{q-1} m_{i-1} = m_{i-q} \bigoplus_{j=0}^{q-1} 2^{q-j} h_j m_{i-q+j} \quad (4.91)$$

where \oplus is the bit-by-bit exclusive-or (XOR) operator:

$$\begin{aligned} 1 \oplus 0 &= 0 \oplus 1 \equiv 1 \\ \text{and} \\ 1 \oplus 1 &= 0 \oplus 0 \equiv 0. \end{aligned} \quad (4.92)$$

For instance: $35 \oplus 27 = 100011 \oplus 011011 = 111000 = 56$.

Step 2: Calculate a set of “direction numbers” by converting m_i into a binary fraction in the base 2 number system. The i -th direction number $v(i)$ for $i = 1, 2, \dots, M$ is given by

$$v(i) = \frac{m_i}{2^i} \quad (4.93)$$

which is always smaller than zero since $m_i < 2^i$. For example, if m_5 is 14 ($14 < 2^5 = 32$), its direction number is given by

$$v(5) = \frac{14}{2^5} = 01110 \gg 5 = 0.01110 = 0 + 0\frac{1}{2} + 1\frac{1}{4} + 1\frac{1}{8} + 1\frac{1}{16} + 0\frac{1}{32},$$

where “ $\gg 5$ ” means “shift the number by 5 digits to the left” (while “ $\ll 5$ ” means “shift the number by 5 digits to the right”). In concrete algorithms, we can thus represent $v(i)$ by the integer $V(i)$:

$$V(i) = 2^{M-i} m_i = m_i \ll M - i, \quad \text{so that} \quad v(i) = \frac{V(i)}{2^M}, \quad (4.94)$$

which in our example would be of the form:

$$V(5) = \underbrace{0111000\dots 0}_{M \text{ times}}.$$

Step 3: Calculate the i -th Sobol number $\theta(i)$ for $i = 0, 1, 2, \dots, N - 1$ using the Antonov and Saleev recursive algorithm:

$$\begin{aligned} \theta(i+1) &= \frac{\Theta(i+1)}{2^M} \\ \Theta(i+1) &= \Theta(i) \oplus V(c) \end{aligned} \quad (4.95)$$

where $\Theta(0) = 0$; $V(c)$ represents the c -th direction number; and c is the rightmost zero-bit in the base 2 expansion of i (e.g., the rightmost zero-bit of $87 = “01010111”$ is the the fourth digit from the right; hence $c=4$).

Example: $N = 1024$, $M = 11$, $(q, p) = (5, 4)$ that is $h_1 = h_2 = h_4 = 0$ and $h_0 = h_3 = h_5 = 1$. Equation (4.91) is: $m_i = 2^5 m_{i-5} \oplus m_{i-5} \oplus 2^2 m_{i-2}$ and we finally set $m_1 = 1$, $m_2 = 3$, $m_3 = 7$, $m_4 = 5$, $m_5 = 29$ and obtain:

i	m_i	$V(i)$	$v(i)$
1	1=00000000001	1024=10000000000	0.5
2	3=00000000011	1536=11000000000	0.75
3	7=00000000111	1792=11100000000	0.875
4	5=00000000101	640=01010000000	0.3125
5	29=00000011101	1856=11101000000	0.90625
6	53=00000110101	1696=11010100000	0.828125
7	23=00000010111	368=00101110000	0.1796875
8	51=00000110011	408=00110011000	0.19921875
9	249=00011111001	996=01111100100	0.486328125
10	881=01101110001	1762=11011100010	0.8603515625
11	1393=10101110001	1393=10101110001	0.68017578125

To generate an n -dimensional Sobol sequence (i.e., n time intervals per price path) beginning with the first dimension, each successive dimension is generated sequentially using a different primitive polynomial. Typically, the procedure begins with the primitive polynomial of lowest degree, and moves to higher degrees as n increases. The steps 1 and 2 are repeated for each of the n dimensions, using a different primitive polynomial for each dimension, while step 3 is executed simultaneously for all dimension, giving an n -dimensional vector.

We have plotted in Figure 4.5.5 some charts of the repartition of the random numbers in different dimensions.

Example: Sobol Random Generator Efficiency


We have implemented a Sobol random generator to simulate HN GARCH(1,1) processes and compared the obtained efficiency with the one of an standard “pseudo” random generator. For the model under consideration we used the parameters published in the paper of Heston and Nandi (1997) evaluated for the SP500, with data ranging from 01/08/92 until 12/30/94.

$$\begin{aligned}
 \alpha &= 1.32e - 6 & (4.96) \\
 \beta &= 0.589 \\
 \gamma &= 422.095 \\
 \omega &= 5.02e - 6 \\
 \sigma_0 &= 9.45e - 3
 \end{aligned}$$

We have estimated the efficiency as function of the dimension (1 day=1 dimension) on the basis of 250 independent prices evaluations for each efficiency estimate with a relative accuracy of 1% for the evaluation of the call prices. We obtain the following results:

Number of Days	Pseudo Random T_{CPU} [s]	MC Efficiency	Sobol T_{CPU} [s]	MC Efficiency	Efficiency Ratio
5	0.4	6.36e-06	0.4	1.13e-07	56.1
20	1.5	7.92e-05	1.5	7.68e-06	10.3
60	4.4	5.49e-04	4.4	1.98e-04	2.77
120	8.9	2.07e-03	8.9	1.42e-03	1.46

The ratios of the two efficiencies have been represented on Figure 4.5.6 which shows that the advantages of a Sobol random generator are very significant until 50 dimensions. For higher dimensions, until 150, it is still advantageous to use a Sobol random generator rather than a uniform one.

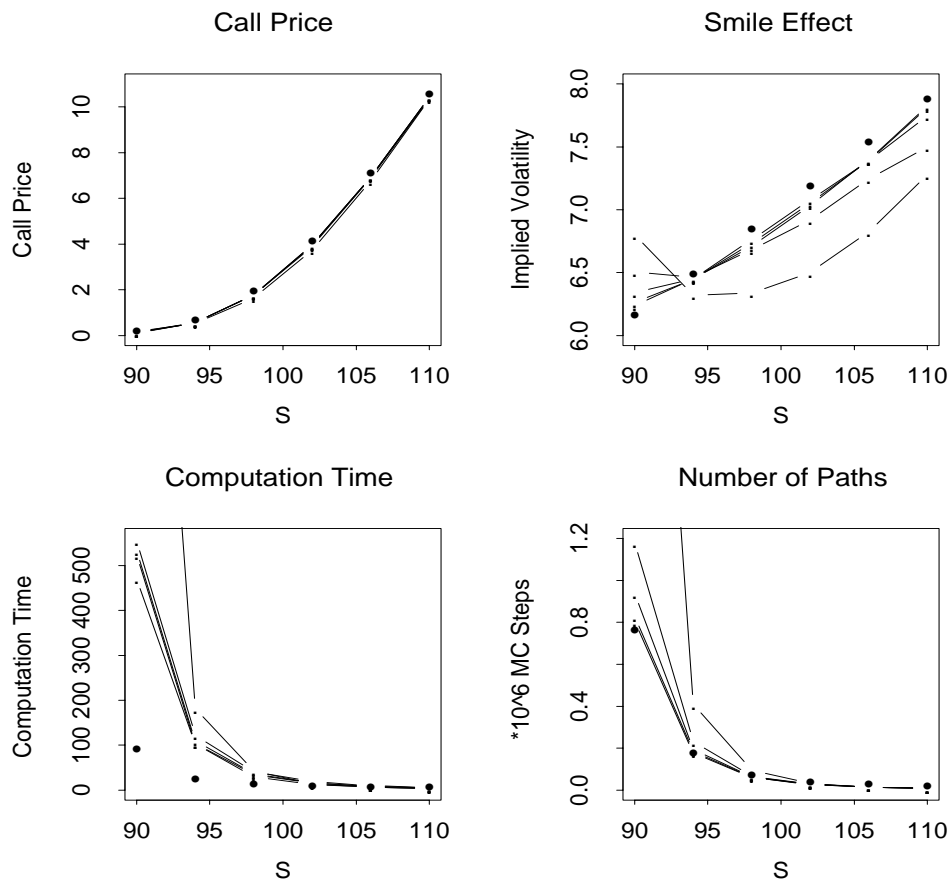


■ Figure 4.5.6: Efficiency of Monte Carlo simulations with a standard uniform random number generator versus Monte Carlo simulations with a Sobol generator as function of the dimension of the domain of integration.

Example: MC GARCH Option Pricing with Leptokurtic Innovations


To observe the influence of a fat tailed distribution on the option pricing, we have used a student- t distribution which is a good candidate allowing for excess kurtosis.

To observe the influence of the excess kurtosis on the option pricing, we have used the max-likelihood estimated parameters from the times-series of the SFE future prices from 1992 to 1999.



■ Figure 4.5.7: GARCH Option Pricing with Leptokurtic Innovations Represented by Points “-.” versus Gaussian Innovations, “*.”. The Prices are calculated with 1% of accuracy. ν go from 3 to 7 and the Prices Converge to the Gaussian Values as ν increases. The computational time is the biggest for ν equal 3.

Figure 4.5.7 and figure 4.5.8 represent the Monte Carlo pricing of a call option with $X = 100$ for different values of the “degree of freedom”, ν , in the Student’s t-distribution. As expected, the smile effect is more important when ν is small. The fat tails increase the price of the option significantly. We have run the Monte Carlo simulations until $S = 90$. Unfortunately as shown in figure 4.5.7, the computational time blows up for extreme out of the money options. To be able to investigate these particular options which are in fact the most difficult to price accurately, more sophisticated Monte Carlo methods like “importance sampling” will become necessary. The problem with “standard” sampling is the fact that with increasing S more and more of the generated paths have zero payoff. A significant variance reduction would be certainly achieved if we made a weighed sampling, generating only the paths that contribute to the integral to be calculated. This will be discussed in the next section.



■ Figure 4.5.8: Yearly Implied Volatility in % as function of $1/\nu$. We see that the Curvature is bigger when ν is small accentuating the Smile effect.

4.5.4 Monte Carlo Importance Sampling

The method of antithetic variates denies efficiently reducing the variance in simulations where prices of deep-out-of-the-money options are to be evaluated. The estimator is still subject to high uncertainties since most of the simulation runs return a value of zero.

The importance sampling technique builds on the observation that an expectation under one probability measure can be expressed as an expectation under another one through the use of a *likelihood ratio*. Its primary focus is to concentrate simulating on sample paths that contribute most to estimate the expected payoff.

Normally, the more a call option becomes out-of-the-money the more simulations are required to maintain a constant variance of the estimator. However, we are free to generate S_T with any other drift μ , provided we weight the result with the appropriate likelihood ratio L_μ . The larger the drift μ , the higher the probability that the option ends in-the-money at maturity and the smaller the number of samples required to estimate the option price, as has been shown by Vázquez-Abad and Dufresne (1998) as well as by Su and Fu (2000) in their studies.

When changing the risk-neutral measure from Q , which involves r , to P , which encompasses μ , the expression for the estimator evaluates to

$$\begin{aligned}
& \mathbb{E}^Q [\max (S_T^r - X, 0)] \\
&= \mathbb{E}^Q [\max(S_T^r(\tilde{X}^r) - X, 0)] \\
&= \mathbb{E}^Q [L_\mu \max(S_T^r(\tilde{X}^r + (\mu - r)\sigma) - X, 0)] \\
&= \mathbb{E}^P [L_\mu \max(S_T^\mu - X, 0)] ,
\end{aligned} \tag{4.97}$$

where the following notations are effective

$$\begin{aligned}
X_k^\mu &= \mu_\sigma \Delta t + \sigma \sqrt{\Delta t} Z_k , \\
\tilde{X}^\mu &= (X_1^\mu, \dots, X_n^\mu) , \\
S_{k\Delta t}^\mu &= S_{(k-1)\Delta t}^\mu e^{X_k^\mu} , \\
\mu_\sigma &= \mu - \frac{\sigma^2}{2} .
\end{aligned} \tag{4.98}$$

In the case of Gaussian innovations, the term for the likelihood ratio calculates to

$$L_\mu = \left(\frac{S_T^\mu}{S_0} \right)^{\frac{r-\mu}{\sigma^2}} e^{\frac{(\mu_\sigma^2 - r_\sigma^2)T}{2\sigma^2}} . \tag{4.99}$$

S_T needs not even be sampled from a log-normal distribution. The only requirement is that the support of the importance sampling measure contains the support of the original measure for the likelihood ratio being well-defined. It is therefore an *absolute continuity* requirement which means that any distribution for S_T whose support includes $(0, \infty)$ is admissible. Hence, the importance sampling technique also holds true when reverting to fat tailed distributions such as the Student-t distribution.

With the use of importance sampling, the simulation procedure becomes similar to the one pricing an at-the-money option and the number of simulation steps needed to attain a certain accuracy threshold remains constant, completely independent of the strike to asset price ratio, see table 4.5.1.

Implementation of the Algorithm

The algorithm we have implemented subsequently encloses the Bollerslev's GARCH(1, 1),

$$\begin{aligned}
\ln S_t &= \ln S_{t-1} + \mu + \lambda \sigma_t^2 + \sigma_t Z_t , \\
\sigma_t^2 &= \omega + \alpha \sigma_{t-1}^2 Z_{t-1}^2 + \beta \sigma_{t-1}^2 .
\end{aligned} \tag{4.100}$$

Duan's N-GARCH(1, 1)

X/S:	1.05	1.15	1.25
Analytic Model			
Black-Scholes Price	1.575	0.205	0.015
HN-GARCH(1, 1) Price	1.196	0.004	0
Antithetic Variates			
Black-Scholes Price	1.583	0.206	0.015
Error	0.51 %	0.46 %	0.57 %
Simulation Steps	80'000	160'000	200'000
Computation Time	7 sec	12 sec	15 sec
HN-GARCH(1, 1) Price	1.201	0.007	0
Error	0.40 %	-	-
Simulation Steps	90'000	200'000	400'000
Computation Time	9 sec	18 sec	42 sec
Antithetic Variates + Importance Sampling			
Black-Scholes Price	1.583	0.206	0.015
Error	0.50 %	0.48 %	0.47 %
Simulation Steps	12'000	12'000	12'000
Computation Time	2 sec	2 sec	2 sec

■ Table 4.6.1: Comparison of computation times in seconds for standard European call options bearing different strike to asset price ratios. The first block (Analytic Models) summarizes exact results, the second block (Antithetic Variates) denotes the calculations performed only taking use of antithetic variates, and the third one (Antithetic Variates + Importance Sampling) sums up the corresponding results with the additional application of importance sampling. Identical parameters concern $S = 100$, $T = 30$ days, $E[\sigma] = 25\%$ and $r = 0\%$. The HN-GARCH(1,1) model parameters are listed in table 4.5.2.

$$\begin{aligned}
\ln S_t &= \ln S_{t-1} + \mu + \lambda \sigma_t^2 + \sigma_t Z_t, \\
\sigma_t^2 &= \omega + \alpha \sigma_{t-1}^2 (Z_{t-1} - \gamma)^2 + \beta \sigma_{t-1}^2.
\end{aligned} \tag{4.101}$$

and Heston and Nandi's HN-GARCH(1, 1) model

$$\begin{aligned}
\ln S_t &= \ln S_{t-1} + \mu + \lambda \sigma_t^2 + \sigma_t Z_t, \\
\sigma_t^2 &= \omega + \alpha (Z_{t-1} - \gamma \sigma_{t-1})^2 + \beta \sigma_{t-1}^2.
\end{aligned} \tag{4.102}$$

as well as the Black and Scholes model as a limiting case. The Black-Scholes and Heston-Nandi models are attractive inasmuch they offer analytical solutions which allow to compare the numerical with the exact results. Contrary, the GARCH and N-GARCH models are only accessible via simulation. However, they impose less stringent constrictions on the paths' properties and can be tailored to fit observed data more accurately.

Virtually all types of options are accessible via the simulation program as it calculates the entire path of the underlying. Currently, our algorithm masters plain vanilla call and put options, and for example arithmetic and geometric Asian options as well as Lookback and Digital options. The library contains different random generators for Gaussian and Student-t distributed innovations. Grace to the modular concept of the software package, further pricing models as well as other distributions are easily addable to the existing ones.

Variance reduction is standardly performed by the use of antithetic variates and is applied on all models. When pricing options bearing a negative moneyness, importance sampling is additionally activated.

As an example for standard normal distributed innovations that drive a Black-Scholes model, the process to generate the paths is specified by

$$S_T^{(i)} = S_0 e^{(r - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}Z^{(i)}} , \quad (4.103)$$

where

$$Z_k^{(i)} \sim \mathcal{N}(0, 1) \quad k = 1, \dots, n \quad (4.104)$$

and $n = T/\Delta t$ denotes the number of steps in each path.

With the Monte Carlo estimate for $\max(S_T - X, 0)$, the pricing formula (??) for a European call figures out to

$$\hat{c} = e^{-rT} \frac{1}{N} \sum_{i=1}^N [\max(S_T^{(i)} - X, 0)] , \quad (4.105)$$

where the estimate is replaced by the average over the possible outcomes minus the strike price.

For the determination of the option's delta we enlisted the assistance of pathwise derivatives based on *infinitesimal perturbation analysis* which result in unbiased estimates under the appropriate conditions. The advantage of this approach lies in the fact that the derivatives are computed from a single simulation without the need for perturbed re-simulations.

Pricing GARCH Options

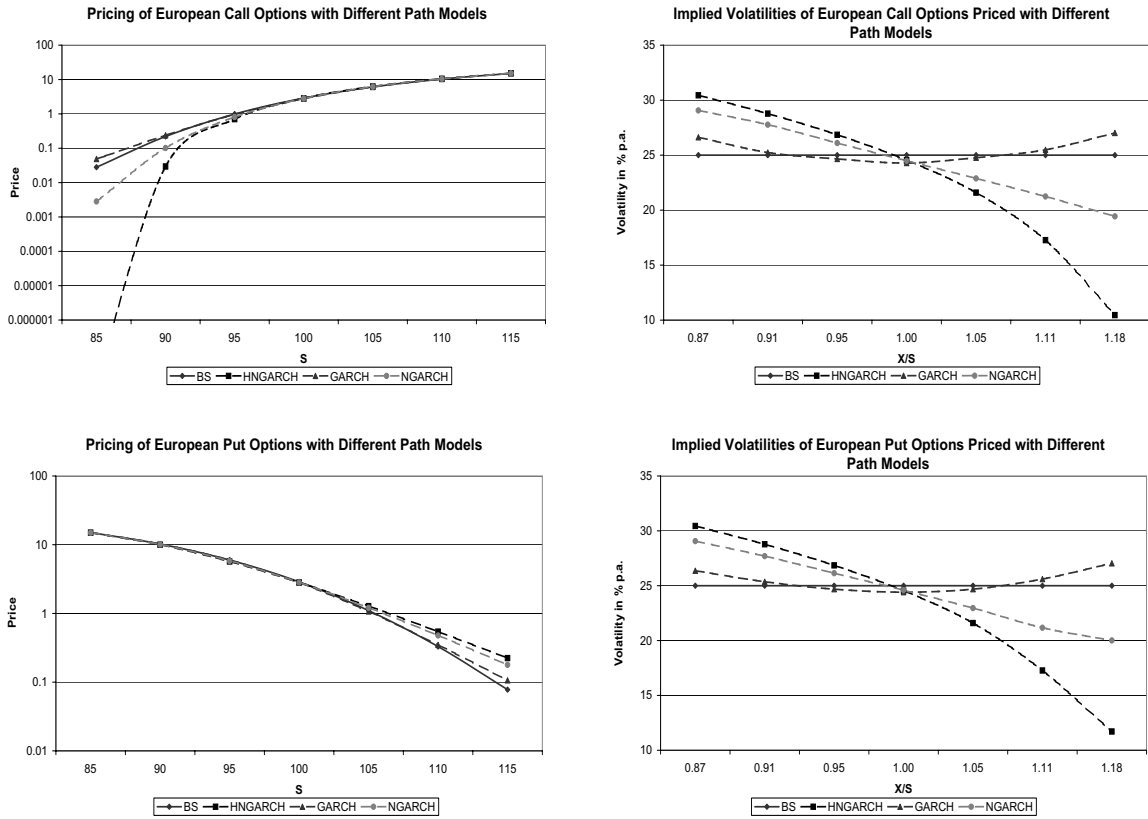
The sector of path-dependent options combined with the appliance of the stochastic volatility models reflects the domain where the Monte Carlo simulation reveals its full potency.

Grace to the programs efficiency, we were able to investigate in the BS, GARCH, N-GARCH and HN-GARCH models in further detail and to examine the consequences on prices when varying the models' different structure parameters. The synopsis of our work is illustrated in figure 4.6.2 where the prices and implied volatilities of the different models are mutually opposed for European calls and puts respectively.

The parameters for the different pricing models are summarized in table 4.6.2 and are either estimated from market data or adopted from published work, see Heston (1993) and Hsieh and Ritchken (2000).

The Influence of Fat-tailed Innovations

The appliance of heavy-tailed distributions generally increases the contracts' values across all models and strike prices as can be verified in figure 4.6.3. It also discloses a notable result by showing that the introduction of fat-tailed innovations alone can produce the characteristic volatility smile observed in real market prices. Remarkably, this property emerges independently on the existence of any correlation between the individual log-returns as is understood by the curved line of the Student-t Black-Scholes prices. Likewise, the curvature of the smile derived from GARCH(1, 1) prices accentuates. The model already embodies a less prominent smile with



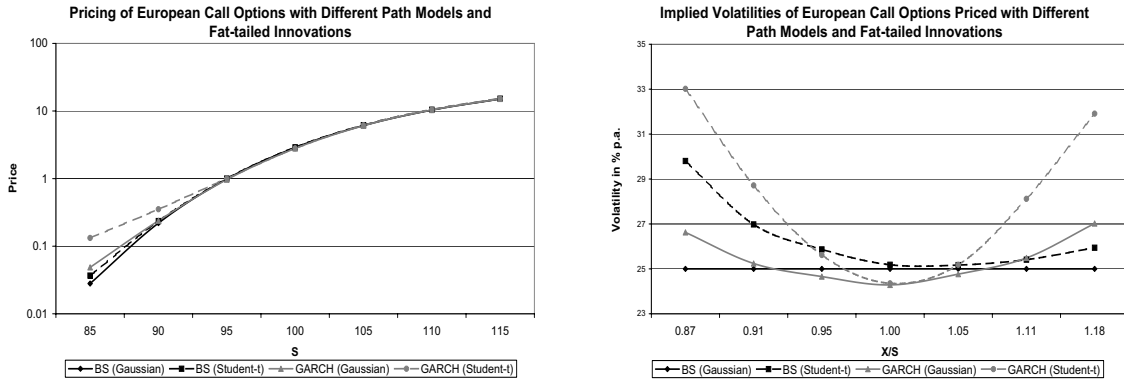
■ Figure 4.6.2: Summary of the results for European call (top row) and put (bottom row) options calculated with the Black-Scholes, GARCH(1,1), N-GARCH(1,1) and HN-GARCH(1,1) models. The left column displays the respective prices, facing their Black-Scholes implied volatilities on the right side. Alike model parameters are $X = 100$, $T = 21$ days, $E[\sigma] = 25\%$ and $r = 0\%$.

Model	ω	α	β	γ	λ
HN-GARCH(1, 1)	$4.939 \cdot 10^{-6}$	$1.579 \cdot 10^{-6}$	$9.062 \cdot 10^{-9}$	785.3	-0.5
GARCH(1, 1)	$4.960 \cdot 10^{-7}$	$7.300 \cdot 10^{-2}$	0.925	0	-0.5
N-GARCH(1, 1)	$8.991 \cdot 10^{-6}$	$6.750 \cdot 10^{-3}$	0.525	8	-0.5


■ Table 4.6.2: Summary of the risk-neutral parameters for the HN-GARCH(1,1), GARCH(1,1) and N-GARCH(1,1) models used to calculate the data displayed in figure 4.6.1, 4.6.2, and 4.6.3.

GARCH(1, 1)	ω	α	β	γ	λ	$\alpha + \beta$
Low Persistence	$6.017 \cdot 10^{-5}$	$7.564 \cdot 10^{-1}$	0.001	0	-0.5	0.757
Medium Persistence	$2.456 \cdot 10^{-5}$	$5.010 \cdot 10^{-1}$	0.400	0	-0.5	0.901
High Persistence	$4.960 \cdot 10^{-7}$	$7.300 \cdot 10^{-2}$	0.925	0	-0.5	0.998

■ Table 4.6.3: Comparison of the risk-neutral parameters of the three GARCH(1,1) models with different levels of persistence, used to calculate the data displayed in figure 4.6.4.



■ Figure 4.6.3: Repercussions of fat tailed distributions on the prices and implied volatilities of European call options calculated by simulations with the Black-Scholes and GARCH(1,1) models. The results obtained with Gaussian innovations are compared to the respective ones where the innovations were drawn from a Student-t distribution with 5 degrees of freedom. Alike model parameters are $X = 100$, $T = 21$ days, $E[\sigma] = 25\%$ and $r = 0\%$.



■ Figure 4.6.4: Effects of different values for the persistence in similar GARCH(1,1) models on standard European call options. The left chart depicts the prices generated by the three models with the corresponding Black-Scholes values, the right side illustrates the implied volatilities. All models yield an equal value for the unconditional variance and the kurtosis. Coinciding parameters are $X = 100$, $T = 21$ days, $E[\sigma] = 25\%$ and $r = 0\%$.

white noise random numbers due to the allowance for an unlimited kurtosis which raises the occurrence of more extreme deviations from the median log-return.

The Role of the Persistence

We performed a thorough investigation in various GARCH(1,1) processes which differ in their structure parameters but exhibit the same macroscopic qualities such as an equal unconditional variance and kurtosis. We chose three models bearing three different levels of persistence (refer to table 4.6.3). Its last column reproduces the persistence $\alpha + \beta$ which controls the decay of random shocks to the conditional variance. The closer to one the longer influence the effects of past volatility fluctuations the future evolution.

The summary is illustrated in figure 4.6.4. The left image shows the prices yielded by the three different models along with the Black-Scholes results. More profound intelligence is reaped from

Model	Plain Vanilla	Geometric Asian	Lookback	Digital
Black Scholes				
21 days	3 sec	4 sec	4 sec	4 sec
45 days	6 sec	8 sec	8 sec	6 sec
90 days	10 sec	13 sec	13 sec	11 sec
HN-GARCH(1, 1)				
21 days	5 sec	6 sec	6 sec	5 sec
45 days	7 sec	9 sec	9 sec	8 sec
90 days	12 sec	16 sec	16 sec	13 sec
N-GARCH(1, 1)				
21 days	4 sec	5 sec	5 sec	5 sec
45 days	7 sec	8 sec	8 sec	8 sec
90 days	11 sec	14 sec	14 sec	12 sec

■ Table 4.6.4: Summary of computation times in seconds for different path models and option types with maturities of 21, 45 and 90 trading days. All computations comprised 50'000 steps and yielded an error to the analytic solution (where available) of less than 0.5 % for at-the-money calls.

the right hand chart which depicts the corresponding implied volatilities. When reverting to the GARCH(1, 1) process

$$Y_t = \lambda\sigma_t^2 + \sigma_t Z_t, \quad (4.106)$$

$$\sigma_t^2 = \omega + \alpha\sigma_{t-1}^2 Z_{t-1}^2 + \beta\sigma_{t-1}^2, \quad (4.107)$$

it is recognized that α controls the magnitude of the innovations' driven random effects, while β governs the decay rate of past events in the variance's evolution in equation (4.107).

Hence, the model with the highest level of persistence reveals the least pronounced smile due to the diminished effects of random shocks on the volatility because of the low value for α . Additionally, the high reading for β induces long-lasting effects by past shocks, therefore enabling a fairly appropriate forecast of future volatility schemes. The combination of these two factors constitutes, that the uncertainty about the forthcoming progression decreases relatively to the other two models.

The more distinctive curvature of the model with the lowest persistence traces back to the fact of a minuscule value for β which causes the influence of the variance's past evolution to vanish. Contrary, the higher reading for α enhances the extension of the individual and unpredictable changes in volatility. Consequently, prices for out- and in-the-money options are raised and hence the volatility smile accentuates.

However, the smile of the model with the intermediate value for the persistence is even more arcuated as in the model with the lowest persistence. This behavior shows that the option prices sensitively depend on the peculiar combination of the model's parameters. In the example on hand, the modestly high value for α still seems to enable ample fluctuations in order to sustain increased prices while the higher figure for β did not yet reach a level where prediction of future volatility formation becomes more coherent. Far from it, it is more likely to provoke a strong enough clustering of extreme events raising the probability of large price swings in short time spans. Out- and in-the-money options are the primary beneficiaries of this constitution and consequently the smile's curvature advances even more.

The simulation's efficiency allowed us to investigate in the repercussions of fat-tailed innovations on option prices. We demonstrated that the renowned smile is at least partly invoked by the introduction of a heavy-tailed distribution alone.

Furthermore, we studied the implications of the GARCH(1,1) model's different structure parameters on option prices and their implied volatilities. We have shown that the dynamics are highly determined by the delicate interaction of the two parameters α and β which control the variance's evolution and hence exert distinct alterations in the options' prices and their implied volatilities.

4.6 The fOptions Library

4.6.1 Summary of Splus Functions

The following section gives an overview over the Splus functions available in the fOptions Library. The programs are grouped by their functionalities. A short description follows each Splus function name.

Pricing Generalized Black-Scholes Options

GBlackScholes	The generalized Black Scholes (GBS) model
GDelta	Delta for the GBS model
GGamma	Gamma for the GBS model
GVega	Vega for the GBS model
GTheta	Theta for the GBS model
GRho	Rho for the GBS model
GCarryOfCost	Carry-of-Cost sensitivity for the GBS model
BAWAmericanApproximation	Barone-Adesi and Whaley American option approximation
BSAmericanApproximation	Bjerk Sund and Stensland American option approximation

Pricing Options with Binomial and Trinomial Trees

CRRBinomial	Standard binomial tree model
BarrierBinomial	Barrier options in binomial trees
Trinomial	Trinomial tree model
ThreeDimensionalBinomial	3-dimensional binomial trees on 2 assets
ImpliedTrinomial	Implied trinomial tree model

Pricing Exotic Options

ExecutiveOption	...
ForwardStartOption	
TimeSwitchOption	
SimpleChooserOption	
ComplexChooserOption	
OptionsOnOptions	
ExtendibleWriterOption	
TwoAssetCorrelation	
EuropeanExchangeOption	

AmericanExchangeOption
 ExchangeExchangeOption
 OnTheMaxMinOption
 SpreadApproximation
 FloatingStrikeLookbackOption
 FixedStrikeLookbackOption
 PartialFloatLBOption
 PartialFixedLBOption
 ExtremeSpreadOption
 StandardBarrierOption
 DoubleBarrierOption
 PartialTimeBarrierOption
 TwoAssetBarrierOption
 PartialTimeTwoAssetBarrierOption
 LookBarrierOption
 SoftBarrierOption
 GapOptionOption
 CashOrNothingOption
 TwoAssetCashOrNothingOption
 AssetOrNothingOption
 SuperShareOption
 BinaryBarrierOption
 GeometricAverageRateOption
 TurnbullWakemanAsianOption
 LevyAsianOption
 FourEquOptInDomCurOption
 QuantoOption
 EquityLinkedFXOption
 TakeoverFXOption

Pricing the Heston-Nandi GARCH(1,1) Option

hngarch.mle	Max-Log-Lokelihood estimation of structure parameters
hngarch.sim	Simulation of time series
hngarch.cvi	Re-evaluation of coniditional variances and innovations
hngarch.mom	Moments of the unconditional log-return distribution
hngarch.momdiff	Partial derivatives of the moments
hngarch.cop	Call and
hngarch.pop	put prices of HN-Garch(1,1) options
hngarch.cogreeks	Greeks of the call and
hngarch.pogreeks	put options
hngarch.coparms	Partial derivatives of call and put option prices
hngarch.poparms	with respect to the structure parameters
hngarc.rms	Root mean square error estimates of structure parameters
	from real option prices

Pricing Options with MC Simulations

...

4.6.2 List of Datasets

...

4.6.3 List of Splus Examples

<code>xmpCRRBinomial</code>	Standard binomial tree model
<code>xmpBarrierBinomial</code>	Barrier options in binomial trees
<code>xmpTrinomial</code>	Trinomial tree model
<code>xmpThreeDimensionalBinomial</code>	3-dimensional binomial trees on 2 assets
<code>xmpImpliedTrinomial</code>	Implied trinomial tree model
...	

4.6.4 Software Packages

Haug's Excel Spreadsheets

The book of Haug (1997) on option pricing formulas includes Excel spreadsheets to evaluate option prices and greeks for a series of plain vanilla and exotic options. The spreadsheets also include programs for option pricing with binomial and trinomial trees. The original Excel spreadsheet programs were the basis to the Splus functions written by D. Würtz.

Heston-Nandi Options Package

The “Heston-Nandi Options Package” includes Splus functions to evaluate option prices and greeks, to simulate the time series process, and to estimate the structure parameters for HN-Garch(1,1) options. The functions were written by Diethelm Würtz and Reto Angliker. Some of the Splus functions are calling Fortran routines. `DUMINF` - a BFGS optimization routine to perform the log-likelihood estimation of the structure parameters, `DQAGI` - a semi-infinite integration routine to evaluate the integrals in the HN option pricing formulas, `DFZERO` - a bisection/secant algorithm in order to search for the zero to solve for the Black-Scholes implied volatility, `DNLS1E` - a Levenberg-Marquard Rroutine to perform the root mean square error estimation of the structure parameters.

Monte Carlo Package

The “Monte Carlo Options Package” includes Splus functions to simulate option prices and greeks based on several variance reduction methods like antithetic variates, low discrepancy sequences and importance sampling. The functions were written by Diethelm Würtz and Beat Bannwart. Some of the functions are calling C or Fortran routines written by the authors and others. The functions are included to the `fOptions` library.

Bibliography

- [1] ... not yet complete
- [2] Abken P.A., and Nandi S. (1996), *Options and Volatility*, Economic Review 12/1996, pp. 21-35.
- [3] Alexander C. (1996), *Risk Management and Analysis*, John Wiley and Sons Ltd, New York.
- [4] Ameur H.B., L'Ecuyer P., Lemieux C. (1999), *Variance Reduction of Monte Carlo Estimators for Stochastic Volatility Models in Finance*, in: Proceedings of the 1999 Winter Simulation Conference.
- [5] Amin K., Ng V., (1993) *ARCH Processes and Option Valuation*, Working Paper, University of Michigan.
- [6] Angliker R. (2000), *Analysis of the Discrete Time Heston and Nandi GARCH Option Pricing Model*, Diploma Thesis, Institute for Theoretical Physics, ETH Zürich.
- [7] Antonov I.A., Saleev V.M. (1979), *An Economic Method of Computing LP_T -sequences*, USSR Comput. Math. Phys. 19, pp. 252-256.
- [8] Armin J., Ng V. (1994), *A comparison of Predictable Volatility Models using Option Data*, Preprint, University of Michigan.
- [9] Baaquie B.E. (1997), *A Path Integral Approach to Option Pricing with Stochastic Volatility: Some Exact Results*, Journale Phys. I France 1997, pp. 1733-1753.
- [10] Bakshi G., Cao C. Chen, Z. (1997), *Empirical Performance of Alternative Option Pricing Models*, Journal of Finance 52, pp. 799-826.
- [11] Bannwart B. (2001) *Efficient Monte Carlo Simulations for Options Pricing*, Diploma Thesis, Institutue for Theoretical Physics ETH Zürich", 83 pages.
- [12] Barlow R.E., Proschan F. (1975), *Statistical Theory of Reliability and Life Testing*, Holt Reinhart and Winston, New York.
- [13] Bartter B., Rendleman R. (1982), *The Pricing of Options on Debt securities*, Journal of Financial and Quantitative Analysis 17, pp. 75-100.
- [14] Bates, D. (1995), *Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options*", Review of Financial Studies 9, pp. 69-107.
- [15] Black F. (1976), *Studies of Stock Price Volatility Changes*", in: Proc. of the 1976 Meetings of the Business and Economic Statistic Section, American Statistical Association, pp. 177-181.
- [16] Black F., Scholes, M. (1973), *The Pricing of Options and Corporate Liabilities*, Journal of Political Economy 81, pp. 637-654.

- [17] Bollerslev T. (1986), *Generalized Autoregressive Conditional Heteroskedasticity*, Journal of Econometrics 31, pp. 307-327.
- [18] Bollerslev T., Chou R.Y., Kroner K.F. (1992), *ARCH Modeling in Finance: A Review of the Theory and Empirical Evidence*, Journal of Econometrics 5, pp. 5-59.
- [19] Bonvin C. (1999), *Low-discrepancy Monte Carlo Simulations of GARCH Processes: Pricing of European Options*, Diploma Thesis, Institute for Theoretical Physics, ETH Zürich.
- [20] Boyle, P. (1977), *Options: A Monte Carlo Approach*, Journal of Financial Economics 4, pp. 323-338.
- [21] Boyle P., Broadie M., Glasserman P. (1995), *Recent Advances in Simulation for Security Pricing*, in: Proceedings of the 1995 Winter Simulation Conference.
- [22] Boyle P., Broadie M., Glasserman P. (1997), *Monte Carlo Methods for Security Pricing*, The Journal of Economic Dynamics and Control 21, pp. 1257-1321.
- [23] Bratley P., and Fox B.L. (1992), *Implementation and Tests of Low-Discrepancy Sequences*, ACM Transactions on Modeling and Computer Simulation 2, PP. 195-213.
- [24] Brennan M. (1979), *The Pricing of Contingent Claims in Discrete Time Models*, Journal of Finance 34, PP. 53-68.
- [25] Broadie, M., Glasserman, P. (1993), *Estimating Security Price Derivatives Using Simulation*, Management Science 42, pp. 269-285.
- [26] Campbell J.Y., Andrew Z.W.L., MacKinlay A.C. (1997), *The Econometric of Financial Markets*, Princeton University Press, New Jersey.
- [27] Charnes J.M. (2000), *Using Simulation for Option Pricing*, in: Proceedings of the 2000 Winter Simulation Conference.
- [28] Christie A. (1982), *The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects*, Journal of Financial Economics 10, 407-432.
- [29] Cox J., Ross S. (1976), *The Valuation of Options for Alternative Stochastic Processes*, Journal of Financial Economics 3, pp. 145-166.
- [30] Cox J., Ingersoll J., Ross S. (1985), *A Theory of the Term Structure of Interest Rates*, Econometrica 53, pp. 385-407.
- [31] Cox J.C., Ross S.A., Rubinstein M. (1979), *Option Pricing: A Simplified Approach*, Journal of Financial Economics 7, pp. 229-263.
- [32] Duan, J.C. (1995), *The GARCH Option Pricing Model*, Mathematical Finance 5, pp. 13-32.
- [33] Duan J.C. (2000), *GARCH and Stochastic Volatility Option Pricing*, Mathematical Finance 3/2000, pp. 1-46.
- [34] Eberlein E., Prause, K. (2000), *The Generalized Hyperbolic Model: Financial Derivatives and Risk Measures*, in: Mathematical Finance - Bachelier Congress 2000, Springer Verlag, New York.
- [35] Eberlein E. (2001), *Recent Advances in more Realistic Risk Management: The Hyperbolic Model*, in: Mastering Risk 2, Prentice Hall - Financial Times.
- [36] L'Ecuyer P. (1994), *Efficiency Improvement and Variance Reduction* Proc. of the 1994 Winter Simulation Conference, pp. 122-132.
- [37] Engle R.F. (1982), *Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation*, Econometrica 50/4, pp. 987-1006.

- [38] Engle R., Mustafa C. (1992), *Implied ARCH Models from Option Prices*, Journal of Econometrics 52, pp. 289-311.
- [39] Exner A. (2000), *Importance Sampling: Theorie und Anwendungen in der Optionsbewertung*, Diplomarbeit, Technische Universität Wien.
- [40] Feller W. (1966), *An Introduction to Probability Theory and Its Applications*, Volume 2, ed. 2, Wiley and Sons, New York.
- [41] Frey R. (1997), *Derivative Asset Analysis in Models with Level Dependent and Stochastic Volatility*, CWI Quaterly 10, pp. 1-34.
- [42] Frey R., Michaud P.K. (1997), *The Effect of GARCH-type Volatilities on Prices and Payoff-Distributions of Derivative Assets - A Simulation Study*, Preprint, ETH Zürich.
- [43] Fu C., Su, Y. (2000), *Optimal Importance Sampling in Securities Pricing*, Proc. of the 2000 Winter Simulation Conference.
- [44] Galanti S., Jung A. (1997), *Low-Discrepancy Sequences: Monte Carlo Simulation of Option Prices*, The Journal of Derivatives 1997, pp. 63-83.
- [45] Glasserman P. (1991), *Gradient Estimation via Perturbation Analysis*, Kluwer Academic publishers, Norwell MA.
- [46] Haug E.G. (1997), *The complete Guide to Option Pricing Formulas*, McGraw-Hill, New York.
- [47] Hammersley J.M., Handscomb D.C. (1967), *Monte Carlo Methods*, Methuen, London.
- [48] Heston S.L. (1993), *A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options*, Review of Financial Studies 6, pp. 327-343.
- [49] Heston S.L., Nandi S. (1997), *A Closed-Form GARCH Option Pricing Model*, Federal Reserve Bank of Atlanta.
- [50] Heston S.L., Nandi S. (1999), *A Closed-Form GARCH Option Pricing Model*, Federal Reserve Bank of Atlanta.
- [51] Hsieh, K.C., Ritchken P. (2000), *An Empirical Comparison of GARCH Option Pricing Models*, Preprint, Weatherhead School of Management, Case Western Reserve University, Cleveland.
- [52] Hull J.C. (1993), *Options, Futures, and Other Derivative Securities*, Prentice Hall, New Jersey.
- [53] Ito K. (1951), *On Stochastic Differential Equations*, Memoirs, American Mathematical Society 4, pp. 1-51.
- [54] Kendall M., Stuart, A. (1977), *The Advanced Theory of Statistics*, Volume 1, Macmillan Publishing, New York.
- [55] Knoch, H.J. (1992), *The Pricing of Foreign Currency Options with Stochastic Volatility*, PhD Thesis, Yale School of Organization and Management.
- [56] Lamoureux C., Lasterapes W. (1993), *Forecasting Stock-Return Variance: Toward an Understanding of Stochastic Implied Volatilities*, Review of Financial Studies 6, 293-326.
- [57] Mandelbrot B. (1963), *The Variation of Certain Speculative Prices*, Journal of Business 36, pp. 394-419.
- [58] Melino A., Stuart T. (1990), *The Pricing of Foreign Currency Options With Stochastic Volatility*, Journal of Econometrics 45, pp. 239-265.

- [59] Merton E. C. (1973), *Theory of Rational Option Pricing*, Bell Journal of Economics and Management Science 4, pp. 141-183.
- [60] Nandi S. (1996), *Pricing and Hedging Index Options under Stochastic Volatility: An Empirical Examination*, Working Paper, Federal Reserve Bank of Atlanta.
- [61] Niederreiter H. (1988), *Low Discrepancy and low Dispersion Sequences*, Journal of Number Theory 30, pp. 51-70.
- [62] Niederreiter H. (1992), *New Methods for Pseudorandom Number and Pseudorandom Vector Generation*, Proc. of the 1992 Winter Simulation Conference.
- [63] Niederreiter H., Spanier J. (1992), *Random Number Generation and Quasi-Monte Carlo Methods*, SIAM, Philadelphia
- [64] Press H.W., Teukosky S.A., Vetterling W.T., Flannery B.P. (1998), *Numerical Recipes in C*, University Press, Cambridge.
- [65] Ross S.M., *A First Course in Probability*, Prentice Hall, New Jersey.
- [66] Rubinstein, M. (1976), *The Valuation of Uncertain Income Streams and the Pricing of Options*, Bell Journal of Economics and Management Science 7, pp. 407-425.
- [67] Spanier J., Maize E.H. (1994), *Quasi-random Methods for Estimating Integrals Using Relatively Small Samples*, SIAM Review 36, pp. 18-44.
- [68] Su Y., Fu M.C. (2000), *Accelerated Simulation for Pricing Asian Options*, University of Maryland.
- [69] Vázquez-Abad F.J., Dufresne D. (1998), *Accelerated Simulation for Pricing Asian Options*, Research Paper Number 62/1998, University of Melbourne.
- [70] Würtz D. (2000), *Analyzing and Modeling Financial Market Data*, Unpublished Material, ETH Zürich.