University of Lausanne

HEC - Master of Science in Banking and Finance

Pricing and Hedging Exotic Options with Monte Carlo Simulations*

Authors Augusto Perilla & Diana Oancea

Supervisor Prof. Michael Rockinger

***Acknowledgements**

We would like to thank Prof. Michael Rockinger for his helpful comments and for his continuous support in achieving our work. A special thank to Peng Cheng for useful references and observations.

August, 2003

Abstract

This paper attempts to implement Monte Carlo simulations in order to price and hedge exotic options. Many exotic options have no analytic solutions, either because they are too complex or because the volatility specification is wrong. Consequently, numerical solutions are a necessity. We discuss the advantages and the drawbacks of such a pricing approach for the main exotic options. Given the strong assumptions of the Black-Scholes world, we attempt to relax them and, in particular, we focus on stochastic volatility models. After a review of the literature, we analyze via simulations the impact of stochastic volatility on the valuation of Asian and spread options. Next we construct and evaluate a dynamic hedging strategy for an exchange option under discrete rebalancing, stochastic volatility and transaction costs. We study the effect of each of these market imperfections on the hedge performance. Finally, we shortly discuss possible hedging approaches for various exotic options and compare static and dynamic hedging.

Table of Contents

Executive Summary

Options that are trivial to price (like binary options) are difficult to hedge. Options that are difficult to price (like Asian options) are trivial to hedge. (Howard Savery, Exotic Options Trader)

Over the last years, the size of the exotic options market has expanded considerably. Today a large variety of such instruments is available to investors and they can be used for multiple purposes. Several factors can provide an explanation for the recent success of these instruments. One possibility is their almost unlimited flexibility in the sense that they can be tailored to the specific needs of any investor. It is why exotic options are also called: "special-purpose options" or "customer-tailored options".

Secondly, these options are playing a significant hedging role and, thus, they meet the hedgers' needs in cost effective ways. Corporations have moved away from buying some form of general protection and they are designing strategies to meet specific exposures at a given point in time. These strategies can be based on exotic options which are usually less expensive and more efficient than standard instruments.

Thirdly, exotic options can be used as attractive investments and trading opportunities. As a result, views on the spot evolution, various preferences on time horizons and premium contingency can all be accommodated by exotic patterns. Moreover, it becomes possible to undertake a very leveraged position which would be unattainable in the spot or standard options market.

The main types of exotic options have been priced either numerically or analytically. A major element in the derivation of the prices has been the construction of hedging or replicating portfolios. Thus, these two issues are strongly interrelated and we consider that, for a global view, it is necessary to discuss them both. The approach we adopt for pricing and hedging is based on Monte Carlo simulations and it is implemented in Gauss.

The rest of this paper is organized as follows. In Chapter 1, we shortly justify the choice of Monte Carlo simulations for pricing and hedging purposes. We construct several Gauss programs in order to price path-dependent options (Asian options, Barrier options, Lookback options) and correlation options (Exchange options, Foreign equity options, Quanto options, Spread options). The results are compared with those obtained following a Black-Scholes approach according to Zhang (1998). The advantages and disadvantages of this numerical approach are weighted against each other for each of the above option types.

In Chapter 2, we extensively review the recent developments in the volatility modelling and compare the various approaches. Consequently, we present both deterministic models (implied volatility and post-dependent volatility) and stochastic ones (diffusions and non-diffusions). Concentrating on the Hull & White (1987) stochastic volatility model, we implement it in Gauss and show how the pricing of different instruments will be affected. Specifically, we shall price, by simulations: an European call, an arithmetic Asian option and a Spread option under Black-Scholes and, then, under the Hull & White (1987) stochastic volatility specification.

Chapter 3 treats the very sensitive subject of hedging exotic options. After a review of the existing hedging approaches, we implement a dynamic hedging approach for a European call. As before, the strategy is developed with the help of Monte Carlo simulations and it will take into consideration various market imperfections such as: discrete rebalancing, transaction costs and stochastic volatility. We compare the hedging costs under different rebalancing frequencies, increasing transaction costs or various parameters for the stochastic volatility process.

The results for the hedging cost and its variability for a European call will serve as benchmark of comparison for the hedging of an Exchange option. The implementation of the dynamic strategy for the Exchange option will also be constructed in Gauss and it will account for all the previously mentioned imperfections. Finally, we shall discuss and compare the two main approaches to hedging other exotic options: static and dynamic hedging. The last part presents our conclusions.

Chapter 1

Pricing exotic options: a simulation approach

As surprising as it may seem given their complexity, exotic options actually have a long history. Some of them existed already when the Chicago Board of Options Exchange started its activity in 1973. However, the trading was very thin and it wasn't until the end of the '70s and the beginning of the '80s that exotic options started to raise more interest. Nowadays, the trading volumes are high and the users are varied: from large financial institutions to corporations, from fund managers to private bankers. Most of these deals take place in the OTC market although some of these options have been also listed in exchanges.

As already stated in the introduction of this paper, there is a vast literature on pricing exotic options and, for the most important ones, solutions have been found numerically or analytically. Our pricing approach is based on Monte Carlo simulations and, in a first step, preserves all the main assumptions of a Black-Scholes environment. Before we proceed, we must detail the Monte Carlo approach to pricing options in general.

In recent years, the complexity of simulation methods has increased tremendously, in a continuous search for accuracy and speed. The Monte Carlo method in particular can be applied for a variety of purposes: valuation of securities, estimation of their sensitivities, assessment of the hedging performance, risk analysis, stress testing, etc. The literature on simulations is voluminous, starting with the seminal paper by Boyle (1977) until the recent papers on quasi-Monte Carlo.

The option pricing theory relies on the idea that the price of a derivative security is given by the expected value of its discounted payoffs. The expectation is taken with respect to a risk neutral probability measure. The Monte Carlo approach is an efficient application of this theory as summarized by the followings:

- simulate a path of the underlying asset under the risk neutrality condition, over the desired time horizon;
- discount the payoff corresponding to the path at the risk-free interest rate;
- repeat the procedure for a high number of simulated sample paths;
- average the above discounted cash flows over the number of paths to obtain the option's value.

The Law of Large Numbers guarantees the convergence of these averages to the actual price of the option and the Central Limit Theorem insures that the standard error of the estimate tends to 0 with a rate of convergence of $\frac{1}{\sqrt{N}}$ where N is the number of simulations. This convergence rate is based on the assumption that the random variables are generated with the use of pseudo-random numbers. It is possible to achieve an even higher rate of convergence provided that quasi-random numbers are used.

Overall, the method proves to be flexible and easy to implement or modify. It can deal with extremely complicated or high-dimensional problems. As shown, the rate of convergence does not depend on the dimension of the problem. In the case of exotic options, the dimension tends to be high since the risk sources are various and/or the payoffs depend on several assets or observation times. Another advantage of the simulation approach is the confidence interval that it provides for the estimate. This interval shows how accurate the estimate really is and if more time and effort are needed for additional precision. Last, the current advances in technology have reduced the computation time and have made the method more attractive.

There are, also, several disadvantages to this methodology: very complicated problems may require a very high number of simulations for an acceptable degree of accuracy and this may be rather time-consuming and expensive. However, many variance reduction methods have been proposed such as: antithetic variables, control variables, importance sampling and stratified sampling, moment matching, etc. Another improvement is to use deterministic sequences of numbers instead of pseudo-random ones. This procedure is particularly useful for high-dimension integrals.

Given that Monte Carlo can be easily extended for complex derivatives and complex stochastic processes, we have decided to implement it in order to price exotic options. We shall first describe the main characteristics of the most important classes of exotic options: path-dependent options, correlation options and other popular types that cannot be integrated in one of the previous categories. In each class, we have chosen the most traded and used exotic options for which analytical formulas have already been found. The majority of these formulas, under Black-Scholes assumptions, can be found in Zhang (1998). The prices obtained by simulations will be compared with the ones given by these formulas in order to assess the performance of our numerical approach.

The basis of all the pricing procedures is represented by the construction of price paths for the underlying asset. Accordingly, for a stock, with or without dividends, the formula that helps generate a random path is the following:

$$S_{t+\Delta} = S_t\, e^{\left[\left(r-g-\frac{\sigma^2}{2}\right)\Delta+\sigma\epsilon\sqrt{\Delta}\right]} \tag{1.1}$$

where

S_t: stock price at time t

Δ: time interval between observations

r: risk-free interest rate

σ: volatility of stock prices

$\epsilon \sim N(0,1)$random

Next, we shall detail the numerical pricing approach for path-dependent options, in particular for Asian options, Barrier options and Lookback options. In all payoffs, w=1 for a call option and w=-1 for a put option. The interval for the simulated price is always calculated at 95% level of confidence.

1.1 Path dependent options

As their name shows, path-dependent options are options whose payoffs at exercise or expiry depend on the past history of the underlying asset price as well as on the spot price at that moment. Furthermore, this path-dependence can be either strong, as for Asian options or weak as for Barrier options. Strong path dependence means that we must keep track of an additional variable besides the asset level at every observation and time. For example, in the case of Asian options, this variable is the average to date of the asset values.

1.1.1 Asian options

Asian options have payoffs that depend on the average value of the underlying asset over some period of time before expiry. This average can be defined in multiple ways: it can be either arithmetic or geometric, weighted or unweighted, calculated with the help of continuous or discrete observations. The payoff is the difference between this average and a pre-defined strike price. For the geometric Asian option the payoff is:

$$\text{GeomOp} = \max\left[wG(n) - wK, 0\right] \tag{1.2}$$

where

$$G(n) = \left(\prod_{i=1}^{n} P_i\right)^{\frac{1}{n}}: \text{ the geometric average of the observed prices}$$

$$K: \text{ the strike price}$$

For the arithmetic Asian option we have:

$$\text{ArithOp} = \max\left[wA(n) - wK, 0\right] \tag{1.3}$$

where

$$A(n) = \frac{1}{n}\sum_{i=1}^{n} P_i: \text{ the arithmetic average of the observed prices}$$

$$K: \text{ the strike price}$$

It is possible to have the so-called average strike Asian options whose payoffs are represented by the difference between the last asset price and one of the above averages. We should also add the flexible Asian options that allocate different weights to price observations and which are now being extensively used.

The majority, however, are European style options based on unweighted arithmetic averages of the underlying asset. Asian options are mostly used in commodity and currency markets because they are cheaper hedging alternatives to a string of standard options. Attempts to value Asian options in general have been done by: Kemna & Vorst (1990), Turnbull & Wakeman (1991), Levy (1992) or Zhang (1998).

The main difference between arithmetic averages and geometric ones in the case of options is that the latter are lognormally distributed, while the former are not. It is the reason why geometric Asian options can be relatively easy to price in a Black-Scholes environment. However, the prices of arithmetic average options can only be approximated using the formulas for the corresponding geometric averages. Monte Carlo simulations have been used quite often to price arithmetic Asian options and the geometric average based option is used as a control variate.

We have implemented two procedures in Gauss, one for geometric Asian options and one for arithmetic Asian options, both for call and put. We have used antithetic variables in both programs in order to reduce the variance. The simulation of stock price paths is the one presented in 1.1 and the parameters used were:

- starting value of the stock price: $S_0 = 100$;
- strike price: $K = 100$;
- time to maturity: $T = 180$ days;
- volatility of underlying asset: $\sigma = 20\%$;
- risk free interest rate $r = 7\%$;
- number of observations: $N = 180$, i.e. daily observations;

For the geometric call option, the price obtained by performing 10000 simulations was 4.0153, with a confidence interval of (3.9516, 4.0789) while the price given by the

Black-Scholes type formula of Zhang (1998) was 4.0180. The price of the arithmetic call option with Monte Carlo was 4.1396, with a 95% confidence interval of (4.0746, 4.2045) and the Black-Scloles formula generated price was 4.1173. It can be seen that the degree of accuracy of the Monte Carlo approach is high. Moreover, this numerical approach can be extremely useful when the averaging period is more complicatedly defined, i.e. when the Black-Scholes formulas become hard to implement.

1.1.2 Barrier options

The oldest type of exotic options is barrier options. The payoff of a barrier option is identical to the one of a standard option if the option still exists at maturity and 0 (or a rebate) otherwise. It simply means that the underlying asset price must stay in some predefined region for the option to be exercised. Depending on how this region is defined, there are two main types of barrier options: the "in" or "knock-in" barrier options and the "out" or "knock-out" ones. The former have a payoff identical to a standard call if and only if the price of the underlying asset hits the barrier, while the latter have this payoff if the barrier is not touched during the option's life. It is also relevant how the barrier is hit; if, initially, the price is under the barrier, so the barrier will be hit from below, we have "up knock-in" and "up knock-out" options. Conversely, if the barrier is hit from above, we have "down knock-in" and "down knock-out" options. Of course, one last classification is in call or put options, so that finally, given all these possibilities, we shall have a total number of eight plain vanilla barrier options. Besides vanilla types, there are many other variations, more or less complicated: time-dependent barriers, Asian barriers, dual-barriers, forward-start barriers window barrier options, etc.

Barrier options are cheaper than standard options and they can be used for various purposes, from hedging to speculation. Some of the main contributions in pricing these instruments belong to: Merton(1973), Goldman, Sosin & Shepp (1979) or Rubinstein & Reiner (1991). While analytical solutions have been proposed for the plain vanilla barrier options under log-normality and risk neutrality, similar ones may not exist for more complicate payoffs. So, Monte Carlo becomes a good candidate for pricing these instruments.

We have constructed two programs for pricing a down knock-in and an up knock-out barrier option. For the down-in option, the parameters are the same as above, except that we have a barrier level of 95, a rebate of 1.5 and the strike price can be either above or under the barrier. The value given by 10000 simulations for a down-in call with a barrier of 95 and a strike of 98 is 3.0047, the interval of confidence, (2.9120, 3.0974) and the Black-Scholes type price is 3.3905. For an up-out option, we must change the barrier level above the current price, say 105, and establish a strike above or under this barrier. For example, a up-out put with a strike of 102 has a value, by simulations, of 4.1139, a confidence interval of (4.0158, 4.2119) while the Black-Scholes price is 3.9853. It can be noticed that the prices exhibit less precision and the convergence is rather slow.

1.1.3 Lookback options

Lookback options have payoffs that depend on the realized minimum or maximum of the underlying asset over a specified period of time, prior to expiry. There are several types of such options: floating strike lookback options, fixed strike lookback options, American lookback options, partial lookback options, etc. The payoff of a floating strike lookback call option is the difference between the settlement price and the minimum price achieved by the stock during the observation period:

$$\text{FloatCall} = \max\left[S_T - m_t^T, 0\right] \tag{1.4}$$

where

S_T: stock price at expiration

m_t^T: minimum price observed during the option's life

For a put, the payoff at expiration is the difference between the maximum of this stock price and the settlement price, that is:

$$\text{FloatPut} = \max\left[M_t^T - S_T, 0\right] \tag{1.5}$$

where

S_T: stock price at expiration

M_t^T: maximum price observed during the option's life

The payoff of a fixed strike lookback option is similar to the one of a standard option except that the terminal price is replaced, for a call, with the maximum of the asset's price and for a put, with the minimum. These options are also called: "Call on the maximum" and "Put on the minimum", respectively, and their formulas are:

$$\text{FxCall} = \max\left[M_t^T - K, 0\right] \tag{1.6}$$

$$\text{FxPut} = \max\left[K - m_t^T, 0\right] \tag{1.6'}$$

where

S_T: stock price at expiration

M_t^T: maximum price observed during the option's life

m_t^T: minimum price observed during the option's life

There are many advantages associated with floating strike lookback options, as Goldman, Sosin & Gatto (1979) show it: "(1) the options would guarantee the investor's fantasy of buying at the low and selling at the high, (2) the options would, in some loose intuitive sense, minimize regret, and (3) the options would allow investors with special information on the range (but possibly without special information on the terminal stock price) to directly take advantage of such information". Nevertheless, these advantages are counterbalanced by the high premiums charged for such instruments. Regarding their pricing, the above authors were the first to study the European-style floating strike lookback options. Conze & Viswanathan (1991) derived explicit formulas for fixed strike European lookback options, for partial lookback options and for American ones. The framework is the Black-Scholes one and the prices are derived by discounting the expected payoffs at the risk-free rate.

The two Gauss programs we have developed deal with both floating and fixed strike lookback options. The parameters are maintained, but we must mention the current minimum and/or maximum, i.e. the ones observed for the past period of the option's life, because they must enter the Black-Scholes type formulas. For the floating strike lookback options we obtained the following results: the call price by simulations was 23.0486 and the confidence interval (22.8227, 23.2746); the put price was 17.8926, with a confidence interval of (17.7613, 18.0239). The corresponding Black-Scholes type prices were: 23.1080 for call and 18.0529 for put. For the fixed strike lookback options, the strike was K=100 and the results were: the call on the maximum had a price of 21.2390 and a confidence interval of (21.1557, 21.3223), while the put on the minimum had a price equal to 19.6392, a confidence interval of (19.6146, 19.6638). The Black-Scholes prices were 21.4923 and 19.6685, respectively.

The pricing of lookback or barrier options can be further developed under more complicated specifications for the asset's variance. In particular, it can be assumed that the variance follows a CEV (constant elasticity of variance) process or a mean-reverting process and simulations can be performed for the variance process as well. Pricing options under stochastic volatility will be treated in detail in the next chapter.

1.2 Correlation options

The payoffs of correlation options or multi-asset options are affected by at least two underlying assets of the following categories: stocks, bonds, currencies, commodities, indices, etc. These assets can be extremely different or they can belong to the same asset class. It is easy to infer that correlation among these assets will have a major role in the pricing and hedging of these instruments. The problems raised by correlation can be significant since it is even more unstable than the variance.

1.2.1 Exchange options

Exchange options give their owner the right to exchange one risky asset for another. Practically, at maturity, the value of one asset is paid while the value of the other asset is received. We have chosen to price and, later in the paper, to hedge, this type of exotic option because it is the basis of its entire class. Actually, many types of correlation options can be transferred into exchange options and studied as such.

The payoff of the option to receive the first asset and pay the value of the second is:

$$\text{ExchOp}_{1,2} = \max\left[S_{T1} - S_{T2}, 0\right] \tag{1.7}$$

while the option to pay the second asset and receive the first has the payoff:

$$\text{ExchOp}_{2,1} = \max\left[S_{T2} - S_{T1}, 0\right] \tag{1.8}$$

where

S_{T1}, S_{T2}: stock prices at expiration time

The first to price these options was Margrabe (1978) and the derivation of the relevant formulas were extensions of the Black-Scholes work. Interestingly, the performance incentive fee, the exchange offer (between the securities of two different companies), the stand-by commitment (for example, a put on a forward contract in mortgage notes) or the margin account can all be thought of as exchange options.

The Gauss program performs first a Cholesky decomposition of the correlation matrix between the two assets. This way, the standard normal random variables that it will generate for constructing the two assets' paths will be correlated. The parameters are:

- starting values of the stocks prices: $S_{01} = S_{02} = 100$;
- time to maturity: $T = 180$ days;
- volatilities of underlying assets: $\sigma_{01} = 20\%$ and $\sigma_{02} = 15\%$;
- correlation coefficient between the assets: 50%;
- risk free interest rate $r = 7\%$;
- no dividends: $g_1 = g_2 = 0$;
- number of observations: $N = 180$, i.e. daily observations;

By performing 10000 simulations, the price of exchanging the second asset and receiving the first was 5.1246, with a confidence interval of (4.9952, 5.2540) while the Margrabe formula gave a price of 5.0820. Similarly, we computed the price of the option to exchange the first option and receive the second, i.e. 5.0468 with an interval of confidence of (4.9288, 5.1647) while the price by Margrabe formula was 5.0820.

The analytical formulas are obviously more efficient in this case since they are simple to implement. However, Monte Carlo becomes more useful when the option has more than three underlying assets such as for basket options.

1.2.2 Foreign-equity options

As their name suggests, these options are on foreign equity, with strike price in foreign currency, but whose payoff will be transformed in domestic currency given the exchange rate existing at expiration. Practically, there are no restrictions on the evolution of the exchange rate. The payoff of such an option is given by:

$$\text{FxEquity} = F_T \max\left[wS_T - wK_f, 0\right] \tag{1.9}$$

where

F_T: exchange rate at time T in domestic/foreign

S_T: stock price in foreign currency at expiration

K_f: strike price in foreign currency

The correlation between the underlying asset process and the exchange rate plays an important role and directly influences the option's payoff. These options can be interesting for speculating or for hedging exposures to foreign markets. The globalization of the financial markets has led to an increase in demand for such products and this tendency will probably continue for the years to come.

The Gauss application will proceed to a similar Cholesky decomposition as in the case of the exchange option. Then, according to the pre-established parameters, it will simulate paths for both the underlying asset and the foreign exchange rate.

These parameters are:

- initial value of the stock price: $S_0 = 100$;
- initial value of the exchange rate: $Fx = 0.8992 EUR/USD$;
- volatility of the stock: $\sigma_1 = 15\%$;
- volatility of the exchange rate: $\sigma_2 = 20\%$;
- correlation coefficient between the stock and the exchange rate: 25%;
- domestic (U.S.) risk free interest rate $r = 8\%$;
- foreign (France) risk free interest rate $r = 7\%$;
- time to maturity: $T = 180$ days;
- number of observations: $N = 180$, i.e. daily observations;

Discounting and averaging the future payoffs of the foreign equity option, we obtained, after 10000 simulations, a price in domestic currency of 7.0728 and an interval of confidence of (6.9221, 7.2235); the price computed with the Black-Scholes type formula was 7.0582.

When the correlation coefficient is 0, the foreign equity option can be priced with the traditional Black-Scholes for standard options, then the payoff is multiplied by the spot exchange rate at that date. Otherwise, the formula will be more complicated, including the correlation coefficient. In general, the simulations result is fairly precise and easy to obtain.

1.2.3 Quanto options

The most popular type of currency-translated options is quanto options or "quantity adjusted options" or "guaranteed exchange rate options". They are foreign-equity options, but with a fixed exchange rate; it means that the investor can benefit from the upward potential of his foreign option without having to worry about exchange rate fluctuations. As a result, the function of a Quanto is simply to transform a foreign contingent claim in domestic currency. Allowing to exchange one risky asset for an option, quantos are called second-order correlation options. Their payoff is given by:

$$\text{Quanto} = F \max\left[wS_T - wK_f, 0\right] \tag{1.10}$$

where

F: pre-determined exchange rate in domestic/foreign currency

S_T: stock price in foreign currency at expiration

K_f: strike price in foreign currency

Quanto options are traded mainly in OTC markets, but they have also been listed in the American Stock Exchange from 1992. In particular, quantos are efficient when used for treasury and commodity risk management of a corporation or for the risk management of an equity or fixed income derivatives book, as Ho, Stapleton & Subrahmanyam (1995) prove it. Simple quanto options can be further complicated with Asian or barrier features and used for specific risks or for implementing market views. The first to price and show how to hedge these instruments was Reiner (1992), followed by Dravid, Richardson & Sun (1993) and Toft & Reiner (1997). Duan & Wei (1999) price quantos under GARCH using simulations. Kwok & Wong (1999) go a step further and price "exotic quantos" (joint quantos with or without barrier, Asian single or multi-asset quantos, etc.) in a Black-Scholes framework.

As for the foreign-equity options, Monte Carlo naturally lends itself to pricing these instruments. We shall illustrate this possibility for a plain vanilla quanto option, but, of course, the Gauss procedures can be further complicated to accommodate path-dependent characteristics or multi-asset features. We have used the same parameters as for the previous option, except that the current exchange rate is fixed and will be used to transform the payoff of the foreign option in domestic currency. After 10000 simulations, the price, in domestic currency for a quanto was 6.6014 with a confidence interval of (6.4605, 6.7424) while the price given by formula was 6.4843.

Quantos may seem very similar to the foreign-equity options. However, the differences remain significant: first, the exchange rate is incorporated differently in the two options (fixed for quantos and volatile for foreign-equity options); second, a quanto option price is obtained by discounting its expected payoff in foreign currency at the foreign interest rate while a foreign-equity option price is calculated by discounted the payoff expressed in domestic currency at the domestic interest rate.

1.2.4 Spread options

A spread option is typically written on the difference between two prices, rates or indices. At expiration, this spread will be compared to the pre-determined strike in order to obtain the payoff of the option. They are also correlation options since the linear relationship between the two underlying assets plays a major role. The payoff is:

$$\text{Spread} = \max\left[awS_{T1} + bwS_{T2} - K, 0\right] \tag{1.11}$$

where

- S_{Ti}: underlying assets' prices at time T, $i = 1, 2$
- K: strike price
- a, b: weights of the two assets in the payoff; $a > 0$, $b < 0$; usually, a=1, b=-1

Spread options are very popular and extensively traded in both the OTC markets and exchanges such as New York Mercantile Exchange from 1994. This popularity is explained by their multiple functions: investors / hedgers can use options on the spread between long-term and short-term treasuries; corporations can hedge the risks of their gross profits (for example, in the oil industry, use options on the spread between crude and refined oil prices), etc.

At the beginning, spread options were priced in a simple Black-Scholes framework, considering that the spread itself was an imaginary asset. Garman (1992) replaced this one-factor model with a better one, a two-factor model. There is a whole literature specialized in credit spread options, such as Longstaff & Schwartz (1995), Das & Sundaram (2000), etc. Spread options can be written on more than two assets and such exotic products will be called "multiple spread options". By now, it is almost impossible to find closed-form solutions for pricing such instruments, but Monte Carlo simulations can overcome this issue easily.

In our case, we shall stick to the simple two-factor model used by Zhang (1998) and its results in a Black-Scholes environment for standard spread options. The Gauss procedure adopts the same parameters as for the exchange option, but in addition, there is a strike price K=2. The prices can be calculated for both call and put, and for different weights (and signs) of the two underlying assets in the payoff. Consequently, following 10000 simulations of price paths, the estimate for a put spread option was 6.1967 and the interval of confidence was (6.0672, 6.3262); the corresponding price given by the Black-Scholes type formula was 6.1307. This procedure shows a high level of accuracy and can be easily extended to more complicated payoffs.

Before ending this subchapter on correlation options, a few additional observations are necessary. The majority of pricing models for correlation options are based on the strong assumption that correlation between the underlying assets is constant. However, it is a well-known fact that this is far from reality and such an assumption could generate serious matters for pricing and hedging. There are several methods for estimating correlation coefficients: historical data, implied correlations, GARCH models. Each one has its pros and cons, but the result is the same, namely: one correlation coefficient which will be considered constant and implemented as such. Many studies concentrate in the direction of estimating one value, as above, but it would be probably more interesting to model the correlation as a stochastic process and analyze the changes in pricing and hedging approaches.

1.3 Other exotic options

This category includes some of the most popular exotic options which cannot be classified in one of the previous classes. Specifically, this group comprises: digital options, compound options, chooser options, contingent premium options, hybrid options, etc. We shall shortly describe them, without implementing their pricing, since it is trivial to change the existing Gauss procedures to account for different payoffs.

1.3.1 Digital options

Digital options have as a payoff a pre-specified amount ("cash-or-nothing"), an asset's value ("asset-or-nothing") or the difference between an asset's value and some fixed number ("gap options"). This payoff will be received conditional on the underlying asset passing some threshold which is why they are also called "binary" or "bet options". If the asset involved in the payoff and the underlying asset are one and the same, then digital options are called "ordinary"; if this is not the case, they are more complicated and they become "correlation digital options". Furthermore, they can be combined with other exotic to form even more "exotic" options.

Due to their simple payoffs, the ordinary digitals are sometimes considered as "basic building blocks" for vanilla options. Pricing them under Black-Scholes has been done extensively by Zhang (1998), for all the previously mentioned types and for both European and American ones. A simulation based approach for ordinary digitals would simply mean reproducing the program for one of the path-dependent options (for example, Asian option) and changing the payoff accordingly. Moreover, one of the programs for correlation options can be changed in order to obtain the payoff of the correlation digital option.

1.3.2 Chooser options

Chooser options give their holder the right to decide at some point in time, but before maturity, whether the option will finally be a put or a call. The various names are quite intuitive in this sense: "you-choose" or "as-you-like".

These options are most welcome when a risk exposure is uncertain or when a view on the market has not been well-defined. The pricing formulas can be obtained in a Black-Scholes environment; if the choice time is the current date, the formula will degenerate in the one of a call (or put); if the choice date is close to maturity, the price will tend to the sum of the corresponding vanilla call and put option prices. Anyway, the price will be higher than the one for standard options since there is an advantage of additional time and information.

A numerical approach can be efficient since it will take into account the possible price paths and compute at the choice date both the call and the put prices for a specific path. Based on these put (call) prices, the higher one will be chosen and, finally, the terminal payoff will be discounted as usual.

1.3.3 Contingent premium options

As their name shows , these options have premiums that will only be paid under some conditions. The main types are: “pay-later options” and “CPOs”.

In the case of the former, the option holder pays neither at the beginning, nor if the option ends up out of the money. However, if the option is in-the-money at maturity, a pre-specified premium must be paid. As expected, the option writers want to be compensated for bearing the risk during the option life, so, finally, the premium they receive is higher than the one for the corresponding standard option. The price of such an option is determined in a similar manner to a normal call (put) except that the premium will diminish the payoff at maturity, so it will have to be discounted accordingly. A simulation approach will simply mean a modification of the final payoff while the rest of the procedure will be identical to the one for standard options.

A CPO is defined as the sum of the payoffs of a corresponding vanilla option and a number of “supershares”, i.e. composite binary options that pay a certain amount of cash provided that the price of the asset reaches a certain region. This option can be priced under Black-Scholes, but simulations are also possible and efficient.

In general, Monte Carlo seems suitable for handling complicated payoffs and structures.

Chapter 2

Stochastic volatility models

2.1 Literature review

The Black and Scholes model for pricing derivatives was a major achievement and it is still the most widely used tool for its purpose in the world of finance. Nevertheless, it relies heavily on a number of assumptions that are, to some extent, unrealistic. Among these are: the ability to trade (hedge) continuously, no transaction costs, constant volatility, continuity of stock price process, independent Gaussian returns, etc.

We shall focus first on the issue of "constant volatility". A well-known paradox is the fact that Black-Scholes model is used to derive implied volatilities from observed option prices. If the stock price followed the Black-Scholes model, in an arbitrage-free market, these implied volatilities should be independent of exercise prices and time to maturity and also constant over time. However, these volatilities vary systematically, creating phenomena such as the "smile effect" and the "term structure of volatility". But the underlying assumption of the model is that prices are lognormally distributed with a constant variance. Black and Scholes themselves found that while their main results seemed to be supported, variance was changing over time.

Empirically, there is considerable evidence that variance moves over time and these changes are not totally predictable. For example, a number of researchers such as Blattberg and Gonedes (1974), Castanias (1979) or Clark (1973) have applied different tools for the distribution of returns, from student t distributions to mixtures of normals in order to account for the changing variance.

Pricing and hedging derivatives under a constant volatility lognormal model is the simplest case and it is related to complete markets. A market with stochastic volatility is an incomplete markets case and it has important implications on hedging as it will be shown later. In general, all stochastic models that try to account for a changing variance can be divided into two main categories: deterministic volatility models and stochastic volatility models. We present a schematic summary of these models below:

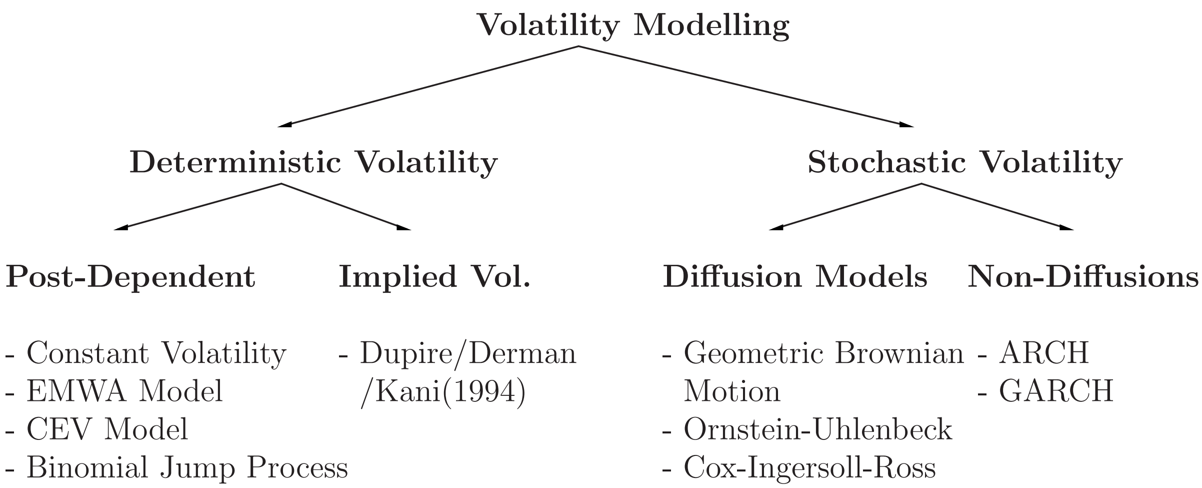

Figure 2.1:

Before presenting the main stochastic volatility models, we take a short look at the deterministic models which assume no stochastic elements in volatility changes. The main issue here is whether the forecast of volatilities should be based on historical or implied volatilities. In general the latter are considered as "better predictors of future volatility" than historical ones, but still most risk management systems use the historical data.

A deterministic volatility model is the Black-Scholes-Merton model itself since it assumes constant volatility of returns over an infinitesimal time period. Stock prices have the Markov property and follow a geometric Brownian motion. The stock prices returns are normally distributed, so the futures stock prices are lognormally distributed. This model is widely used even though it does not fit the real stock return data, as stated previously.

Another possibility of modelling volatility is the so-called Exponential Moving Average of Historical Observations (EWMA) under which the latest observations have the highest weight. Obviously, the difficulties consist in determining the decay factor because the higher it is, the lower the responsiveness of the estimates to new information. The model has the advantage of relatively little data needed and it is used, for example, by J.P.Morgan, with a decay factor of 0.94 and 25 past observations.

The Constant Elasticity of Variance model was developed by Cox and Ross (1976). Volatility is dependent on both the asset price and time and the process followed by stock prices is:

$$\begin{aligned}\frac{dS}{S} &= \mu\,dt + \sigma(S,t)\,dz \\ &= \mu\,dt + \sqrt{\sigma^2 S^{2\rho-2}}\cdot dz \\ &= \mu\,dt + \sigma S^{\rho-1}\,dz\end{aligned} \tag{2.1}$$

It can be shown that the elasticity of the variance is constant. The model is based on the assumption of negative correlation between stock returns and volatility. Two arguments are used to justify this assumption: variance of stock prices when prices are high is not considered dangerous by investors and financial leverage is inversely related to the value of the firm.

In other models, volatility can be just a function of time, $\sigma(t)$. As Hull (1999) explains it: "volatility tends to be an increasing function of maturity when short-dated volatilities are historically low. This is because there is then an expectation that volatilities will increase. Similarly, volatility tends to be a decreasing function of maturity when short-dated volatilities are historically high. This is because there is an expectation that volatilities will decrease". In this case, Black-Scholes still applies, but with a modified volatility, i.e. a time-averaged one:

$$\bar{\sigma}^2 = \frac{1}{T-t}\int_t^T \sigma(s)\,ds \tag{2.2}$$

To have a smile, however, means that volatility must depend on the strike price. As a result, there are many competing ways, both parametric and some nonparametric, to estimate the so-called volatility surface from traded option prices. Such models have been called "implied deterministic volatility models". The advantage of these procedures is that they still assume complete markets. Hence, they give the possibility of deriving hedging strategies and determining option prices which match the observed volatility behavior. However, there may be some problems with the stability of the fits over time, and an extensive empirical study of the stability of these fitted surfaces can be found in Dumas, Fleming and Whaley (1998).

As shown before, there is a class of non-diffusion stochastic volatility specifications that includes ARCH and GARCH models. Autoregressive Conditional Heteroscedasticity or ARCH models were introduced by Engle (1982). They are generalizations of the standard models such as EWMA. Practically, the estimate of variance depends on n previous observations and, eventually, a long-run average variance.

The ARCH model has been successful due to the fact that, while staying simple, it accounts for many features such as fat tails of return distributions or volatility clusterings, i.e. periods with high (low) volatility that tend to increase (decrease) subsequent volatilities. However, it has the disadvantage that negative variances may occur for long lag periods and it is relatively inflexible.

In 1986, Bollerslev proposed the Generalized ARCH or the so-called GARCH which exhibits a longer memory and a more flexible lag structure. One of the main contributions of GARCH is that it recognizes that in practice, volatility tends to be mean-reverting. Moreover, it accounts for the term structure of volatility present in option prices.

In "pure" stochastic volatility models, the asset price satisfies the well-known differential equation :

$$dS_t = \mu S_t \, dt + \sigma_t S_t \, dW_t \tag{2.3}$$

where

$$(\sigma_t)_{t \geqslant 0} : \text{ volatility process}$$

The volatility process can be a simple diffusion, a jump process or a jump-diffusion one. The main difference from the previous deterministic models is that volatility is modeled so that it has a random component of its own. Moreover, the volatility process is not perfectly correlated with the Brownian motion characteristic to the asset prices.

Most papers on this issue follow a similar procedure: first, stochastic volatility option prices should be the solution to a PDE that depends on two state variables: asset price and volatility. However, volatility is neither traded, nor spanned by the existing assets in the economy. So, volatility risk cannot be eliminated through arbitrage arguments and its risk premium enters into the PDE. In general, this premium is assumed to be 0 or some fraction of volatility. Moreover, the volatility process is explicitly modeled as a geometric brownian motion (Hull & White 1987, Johnson & Shanno 1987), an Ornstein-Uhlenbeck process (Scott 1987, Stein & Stein 1991) or a Cox-Ingersoll-Ross one (Heston 1993, Ball & Roma 1994). The PDE is subject to boundary conditions and may lead to a solution for the option's price, but usually this is extremely difficult to find. Nevertheless, a number of authors have applied numerical techniques in order to assess the accuracy of their model.

Hull & White (1987) proposed a model for stochastic volatility that has become a classical reference and a crucial point in the research related to this subject. Variance is assumed to follow a lognormal process, similar to the one traditionally followed by the asset price. We shall provide a full description of this model and apply its associated numerical procedures in the next subchapter.

Another well-known paper which is also conceptually close to the Hull & White approach belongs to Johnson & Shanno (1987). The processes followed by the asset price and the volatility are:

$$dS = \mu S\,dt + \sigma S^{\alpha}\,dz \qquad (\alpha \geqslant 0) \tag{2.4}$$

$$d\sigma = \mu\sigma\,dt + \xi\sigma^{\beta}\,dz_S \qquad (\beta \geqslant 0) \tag{2.4'}$$

The implementation of the model is based on Monte Carlo. The results are as follows: first, the value of out-of-the-money calls increases with the correlation coefficient, while the value of at-the-money options is rather insensitive to this parameter; the implied volatilities increase with exercise price if correlation is positive and decrease with exercise price in the opposite case; the implied volatility tends to be larger for short-term options than for long-term ones. In conclusion, the influence of correlation on option prices seems to be important, but it is difficult to understand why changes in the sign of this correlation may take place.

Scott(1987) uses an independent diffusion process with mean-reversion to model the random changes in implied volatilities from one day to the other. He also proposes a second Ornstein-Uhlenbeck process for the log volatility. The differential equations corresponding to the asset prices and to the volatility are, respectively:

$$dS = \mu S\, dt + \sigma S\, dz_1 \tag{2.5}$$

$$d\sigma = \mu\,(\bar{\sigma} - \sigma)\ dt + \xi S\, dz_2 \tag{2.5$'$}$$

The parameters of the volatility process are estimated with the method of moments and a first order autoregressive process of intra-day volatilities. The starting value of volatility is determined from at-the-money options by OLS, for both the stochastic variance model and for Black-Scholes. Then this value is used to determine the prices of out- or in-the-money options, based on a Monte Carlo procedure. The random variance model performs slightly better than Black-Scholes in explaining actual option prices, and both tend to overprice out-of-the-money options.

In the same line of thought, Stein & Stein (1991) proposed a model according to which the volatility follows an arithmetic Ornstein-Uhlenbeck process. The approach is more general than in the Hull & White model. First, a closed-form exact solution as well as an approximation for the stock price distribution are derived with the help of a Fourier inversion. Then, the parameters are estimated empirically from implied volatilities and used to compute option prices which are compared to Black-Scholes prices. As a result, it is observed that stochastic volatility has an upward influence on all prices and it can account for the smile. Finally, there is a strong link between the parameters of the volatility process and the "fat tails" of the price distribution.

Stein & Stein assume that asset prices are uncorrelated with volatility. Their approach fails to capture the skewness effects that such a correlation implies. Heston(1993) provides a closed-form solution for a European call for the case when the spot asset is correlated with volatility. Furthermore, the model allows for stochastic interest rates and can be used for pricing bond and foreign currency options. The stock price follows the usual diffusion process while the variance is supposed to evolve according to the well-known square-root process used by Cox, Ingersoll and Ross (1985):

$$dS(t) = \mu S\, dt + \sqrt{V(t)} \cdot S\, dz_1(t) \tag{2.6}$$

$$dV(t) = a\,(\bar{v} - v(t))\ dt + \sigma\sqrt{v(t)}\, dz_2(t) \tag{2.6$'$}$$

Correlation between the asset prices and volatility can explain the skewness of the spot returns. If correlation is positive, as stated before, then the stochastic volatility model will also provide higher prices for out-of-the-money options and lower prices for in-the-money ones compared to Black-Scholes. The volatility of the volatility is responsible for the kurtosis and fat tails. One last observation is related to the fact that Black-Scholes and the stochastic volatility models provide almost identical prices for at-the-money options. Since most options trade near-the-money, this may partially explain why Black-Scholes is still the most successful model for option pricing.

Ball & Roma (1994) use a CIR model and compute option prices via integration in Gauss. These prices are compared to the Black-Scholes ones, the Stein & Stein ones and to the Hull & White approximations. Stein & Stein have identified a systematic tendency of the Black-Scholes model to overprice options. In contrast, Ball & Roma prove that stochastic volatility can have both upward and downward effects on option prices as a function of the moneyness of the options. Moreover, it is shown that stochastic volatility can account for the "smile effect".

In conclusion, the stochastic volatility approach:

- models the empirically observed random behaviour of market volatility;
- it allows for skewness due to the correlations between underlying processes;
- it produces more realistic return distributions, including their fat tails;
- it accounts for smile/smirk effect in option prices.

However, in order to be objective, we have to mention the difficulties that such a complex approach may raise: first, volatility is not directly observed. As a result, identifying the correct parameters for a specific volatility process is a most demanding task. There are many potential techniques, mostly relying on implied volatilities and on time-series techniques. Second, the market is incomplete, so derivatives cannot be perfectly hedged with just the underlying asset. Last, there is not a generally accepted model, so a more or less subjective choice must be done among the existing ones.

2.2 The Hull & White model and its applications under Monte Carlo

We shall concentrate now in the description and applications of the most famous stochastic volatility model: Hull & White (1987). The paper produces a power series solution for the case in which the stock price is uncorrelated with the volatility. As stated previously, the process followed by the variance is a lognormal one and it is very similar to the one followed by the security price:

$$dS = \mu S\, dt + \sigma S\, dz_1 \tag{2.7}$$

$$dV = \mu_S V\, dt + \xi V\, dz_2 \tag{2.7'}$$

The drift and the volatility of volatility are independent of price movements, but may be functions of time. The 0 correlation between asset price and volatility is equivalent to assuming a constant volatility of firm value and no leverage. The volatility has no systematic risk which means that its associated risk premium is 0 or that it is uncorrelated with aggregate consumption. Under these specifications, the PDE that a derivative must satisfy is reduced to:

$$\frac{\partial f}{\partial t} + \frac{1}{2}\left[\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} + 2\rho\sigma^3\xi S \frac{\partial^2 f}{\partial S\,\partial V} + \xi^2 V^2 \frac{\partial^2 f}{\partial V^2}\right] - rf = -rs\frac{\partial f}{\partial S} - \mu_S \sigma^2 \frac{\partial f}{\partial V} \tag{2.8}$$

The risk-neutral procedure is applied as usual and they conclude that even if the variance is stochastic, the terminal distribution of the stock price given the mean variance remains lognormal with mean: $rT - \bar{V}\frac{T}{2}$ and variance: $\bar{V}T$. The log-normality is insured only if the stock price and volatility are uncorrelated. Moreover, should the investors be risk averse, the drift of the stock price would be affected by the variance. So, the mean of the terminal stock distribution will depend also on the stochastic variance and log-normality will not necessarily appear.

Under the previous conditions, the option price can be computed using the classical Black-Scholes formula with the average variance replacing the constant one: "the option price is the B-S price integrated over the distribution of the mean volatility":

$$C\left(\bar{V}\right) = S_t N(d_1) - X e^{-r(T-t)} N(d_2) \tag{2.9}$$

where

$$\begin{aligned} d_1 &= \frac{\ln\left(\frac{S_t}{X}\right) + \left(r + \frac{\bar{V}}{2}\right)(T-t)}{\sqrt{\bar{V}\,(T-t)}} \\ d_2 &= d_1 - \sqrt{\bar{V}\,(T-t)} \\ \bar{V} &= \frac{1}{T}\int_0^T \sigma^2(s)\,ds \end{aligned}$$

This pricing formula always holds in a risk-neutral environment provided the 0 correlation between the stock price and the volatility. It is particularly interesting that if, in addition, the volatility is uncorrelated with aggregate consumption, this equation will hold in a risky world as well. If we add the assumptions of 0 drift in the volatility process and small values of the volatility of volatility parameter, then the call price can be approximated by a series expansion that converges rather fast. Examining the second derivative of the option price, it becomes clear that the Black-Scholes price always overprices at-the-money options and underprices deeply out-of- or in-the-money ones. Even though counterintuitive, it seems very plausible that stochastic volatility leads, in some cases, to lower option values than those obtained with a constant volatility model.

Hull & White examine several numerical procedures, particularly Monte Carlo methods, showing that they can be efficiently used for deriving option prices. In addition, the simple lognormal process for variance could be replaced by a mean-reverting one of the form:

$$dV = a\left(\bar{\sigma} - \sigma\right) V\,dt + \xi V\,dz \tag{2.10}$$

where

a: speed of mean-reversion

$\bar{\sigma}$: long-run volatility mean

ξ: volatility of volatility (2.11)

Actually, it is only necessary to simulate the variance process, step by step, then take the arithmetic mean of the variance over a large number of possible paths and use these means in the Black-Scholes formula instead of the constant variance. However, if the correlation between stock price and volatility is nonzero, a more complex approach is necessary [1], involving the simulation of correlated random numbers. Such a modelling is subject to errors due to the sensitivity to parameters. For example, the volatility of the volatility process can be estimated from changes in implied volatilities or from changes in historical variance series. But there is always the danger of mispricing options in the first case, or the lack of data in the second one. Next, the initial value of volatility can lead to significant pricing biases. Finally, the sign of the correlation between asset prices and volatility will determine different pricing biases for the same degree of moneyness. When there is positive correlation, in-the-money options are overpriced by Black-Scholes while out-of-the-money options are underpriced. The effect is reversed for negative correlation. The changes in the correlation sign in different time periods could explain the empirical observations on option prices, but it is rather hard to understand why such changes may occur.

Following the approach suggested by Hull & White, we have implemented a stochastic volatility procedure for pricing a European call option. We assume that stock prices and volatility are uncorrelated. The variance is modelled as a mean-reverting process with the following parameters:

- initial volatility of underlying asset: $\sigma = 15\%$;
- speed of mean reversion: $a = 10$;
- long run volatility mean: $\bar{\sigma} = 15\%$;
- risk free interest rate $r = 0\%$;
- time to maturity: $T = 180$ days;
- number of observations: $N = 180$, i.e. daily observations;

[1] We shall detail these numerical procedures, for a simple call, using Monte Carlo simulations in Gauss, but first we shall focus on the results obtained by Hull & White.

The volatility at each point in time was generated using the stochastic process given by:

$$V_{t+\Delta} = V_t \, e^{\left[a\left(\bar{\sigma}-\sqrt{V_t}\right)-\frac{\xi^2}{2}\right]\Delta+\xi \upsilon \sqrt{\Delta}} \tag{2.12}$$

We generate a vector of N stochastic variances with the help of N independently and normally distributed variables, i.e. υ_1,...,υ_N. At each moment in time, each variance is calculated based on the previous one and a random draw, as specified. These variances are afterwards used to generate a vector of stock prices, according to the formula:

$$S_{t+\Delta} = S_t \, e^{\left[\left(r-\frac{V_t}{2}\right)\Delta+\sqrt{V_t}\epsilon\sqrt{\Delta}\right]} \tag{2.13}$$

The paths for volatility and for the stock prices, respectively, were generated with the help of antithetic variables. As a result, we obtained four price paths and, thus, four possible prices for the call at each simulation run. We call them: P_1, P_2, P_3 and P_4 and they correspond to the following combinations of random draws: (+ϵ, +υ), (+ϵ, -υ), (-ϵ, +υ) and (-ϵ, -υ). In addition, we computed by simulation the prices of the call under constant volatility, but using the same normal draws for ϵ and their antithetic. Thus, we obtained two call prices, Q_1 with ϵ and Q_2 with -ϵ.

Given that our objective is to estimate the bias between the call prices under stochastic volatility and the call prices under constant volatility, we calculate:

$$\begin{gathered} \frac{P_1+P_3-2Q_1}{2} \\ \frac{P_2+P_4-2Q_2}{2} \end{gathered} \tag{2.14}$$

Repeating this procedure 5000 times and averaging these values, we obtained an estimation of the bias. The graph 2.2 shows this bias of prices, in absolute value, for varying values of $\frac{S}{K}$ and volatility of volatility.

We can verify the Hull & White result according to which the Black-Scholes formula overprices at-the-money options and underprices deeply out-of or in-the money ones. In addition as the volatility of volatility increases, this effect becomes more visible. For

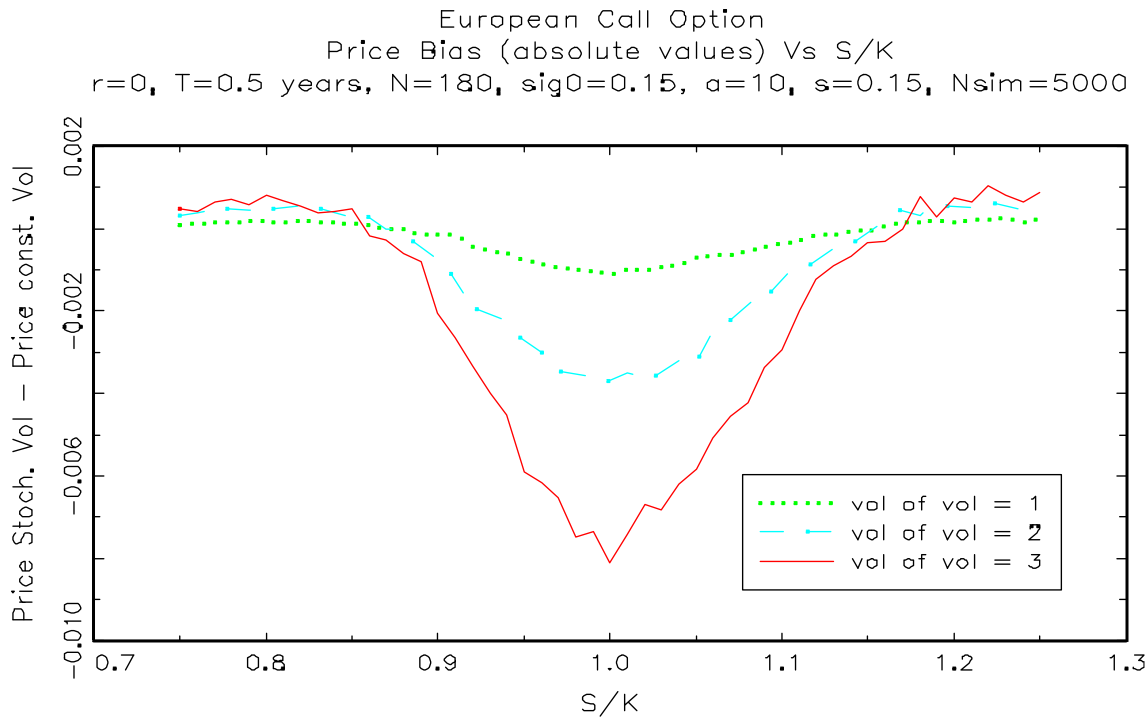

Figure 2.2:

a better understanding of the size of these differences, we chose to represent them as percentages of the corresponding Black-Scholes values (see Figure 2.3).

For example, when the volatility of volatility is 2, the price bias for options slightly out-the- money is of 13%. Moreover, the underpricing effect is bigger for slightly out-the-money options that for in-the-money ones.

The stochastic process for volatility can be correlated with the price process. For this purpose, the corresponding standard normal draws (here, ϵ and υ)will be correlated by means of the familiar by now Cholesky decomposition. Furthermore, our procedure was implemented in Gauss and can be easily adapted for modelling variance with a different process by simply changing the procedure that generates the vector of variances.

2.3 Pricing exotic options with stochastic volatility

A European call presents a significant bias between the price computed under stochastic volatility and the Black-Scholes one. We shall turn our attention to the main subject of this thesis, exotic options. We shall check if this bias holds for them too.

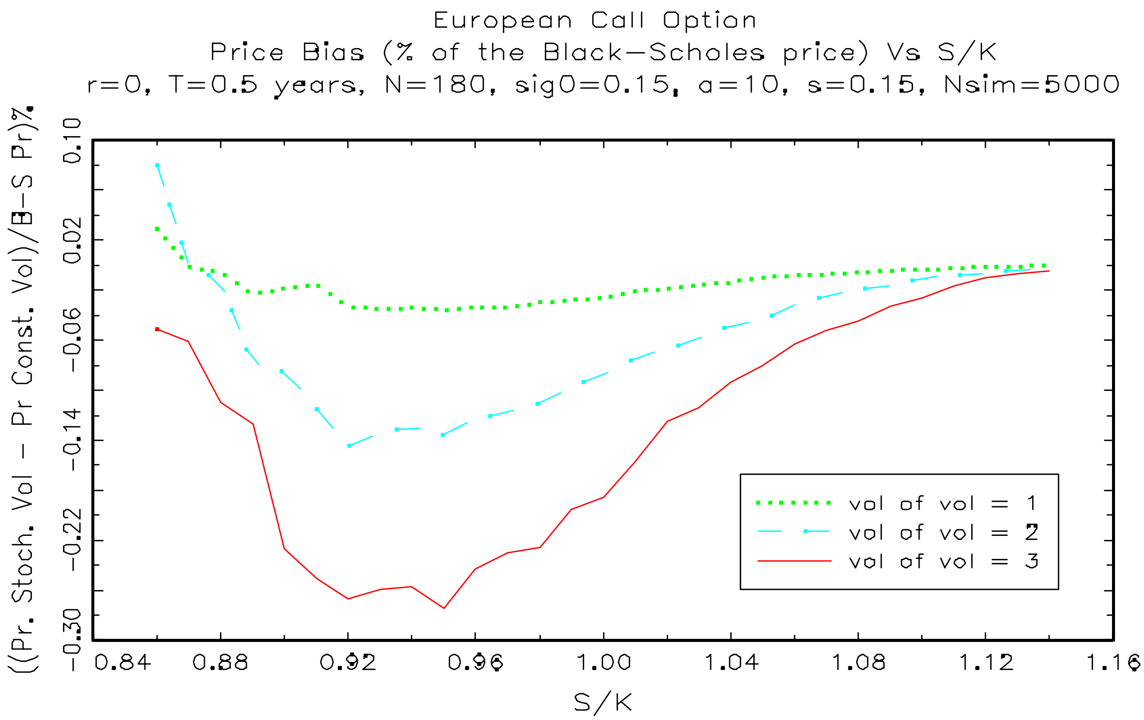

Figure 2.3:

In particular, we chose one path-dependent option, that is an arithmetic Asian one, and a correlation option, i.e. a spread option. For the former, the parameters chosen were:

- initial volatility of underlying asset: $\sigma = 15\%$;
- speed of mean reversion: $a = 10$;
- long run volatility mean: $\bar{\sigma} = 15\%$;
- risk free interest rate $r = 0\%$;
- time to maturity: $T = 180$ days;
- number of observations: $N = 180$, i.e. daily observations;

The simulation procedure is very similar to the one for the European call. The vector of simulated values for variances is used as an input for the price process. The antithetic variables for both the volatility paths and price paths help reduce the variance. In addition, they help obtain four potential prices at each simulation run. The Gauss

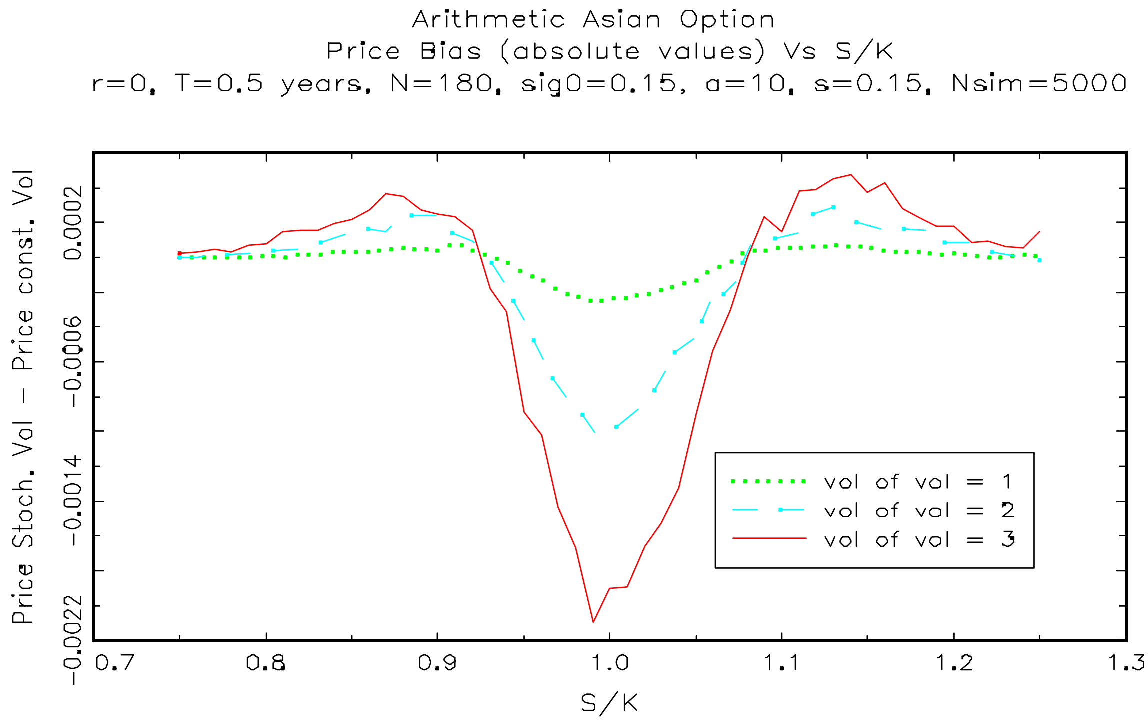

Figure 2.4:

program for pricing Asian options under constant volatility (see Chapter 1) will be run in order to generate the "Black-Scholes" type of prices. The bias is the difference between the price under stochastic volatility and the one under constant volatility. Using the same formulas (formula 2.21), we compute the bias an represent it graphically above for different values of $\frac{S}{K}$ and volatility of volatility.

We observe that, as for a simple call option, the Black-Scholes formula overprices at-the-money options and underprices deeply out-of or in-the money ones. In addition as the volatility of the volatility increases, this effect becomes more pronounced. However, there are certain significant differences. Firstly, the overpricing is concentrated for very near at-the-money options in the case of an Asian call. On the contrary, for a standard call, the overpricing extends on a larger area around the at-the-money point. Secondly, the Asian option exhibits less overpricing in this area than the European call option. Thirdly, it seems that out-of-the money and in-the-money Asian options are more severely underpriced than their counterparts.

In order to have a better image of the previously stated differences, we represent the bias as a percentage of the price obtained under constant volatility in Figure 2.5.

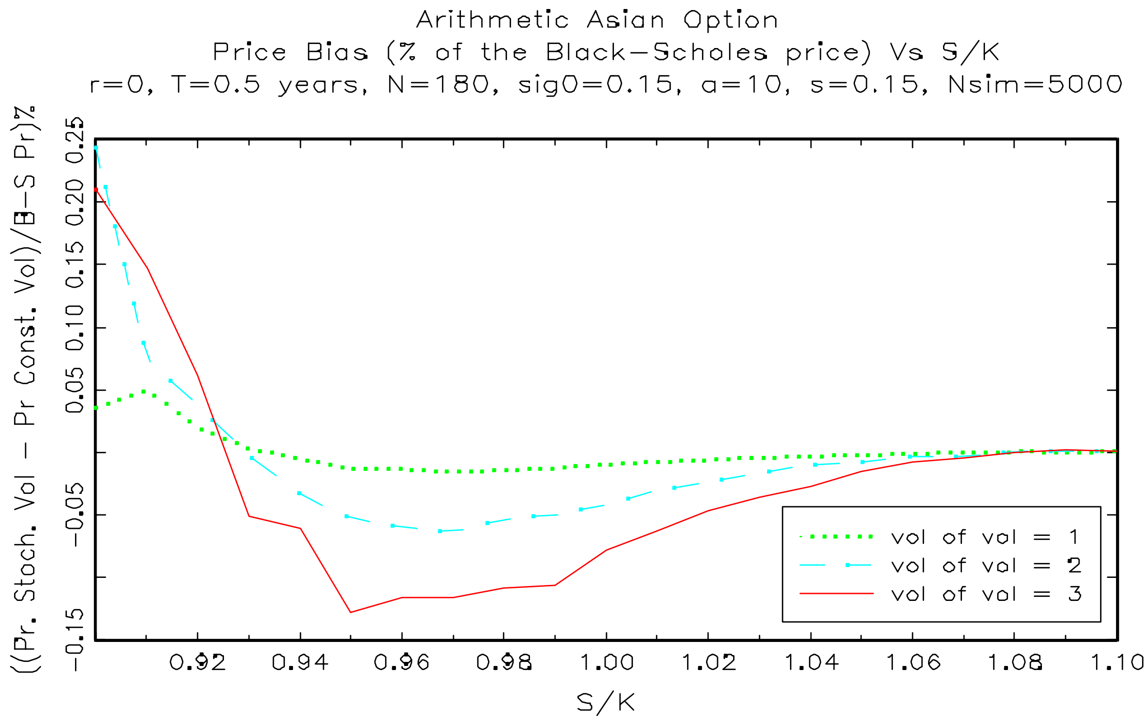

Figure 2.5:

For the particular case of our parameters, we see that the biggest bias is for slightly out-the-money options. The case of deeply out-the-money options, even though it exhibits a huge price bias, is not so important given that the price of the option is almost 0 for these cases.

The second exotic option we chose is a spread one. Ex-ante, we expect the price bias to be more significant given that there are two assets involved in the payoff, each one with its own stochastic volatility process. In particular, we considered that the variances of both assets follow the same type of stochastic process, i.e. a mean-reverting one. The parameters were:

- initial volatilities of the underlying assets: $\sigma_1 = \sigma_2 = 15\%$;
- correlation coefficient between the two assets: $\rho = 0.82$;
- speeds of mean reversion: $a_1 = a_2 = 10$;
- long run volatility means: $\bar{\sigma}_1 = \bar{\sigma}_2 = 20\%$;
- risk free interest rate $r = 0\%$;

- strike price: $K = 0$;
- time to maturity: $T = 180$ days;
- number of observations: $N = 180$, i.e. daily observations;

The procedure for obtaining the spread price bias is more complicated than the previous ones. This is simply due to the fact that for each asset we had to simulate a volatility path; then we used this paths to generate price paths, but including antithetic draws from a standard normal variable. We used the same price paths, but with constant volatility and computed the value under the assumptions of Black-Scholes. For all the alternatives of our antithetic approach, we computed the price biases and averaged them over 1000 simulations.

The simulations were repeated under a large number of possible scenarios, in order to have an accurate view of the effect of stochastic volatility modelling. In particular, provided the above parameters, we changed the initial values of the two assets and we also considered three different possibilities for the volatility of volatility parameter. The results are synthesized in figure 2.5. The striking aspect of this table is the fact that spread options seem to be constantly underpriced by a Black-Scholes type of procedure. However, the result is consistent with our ex-ante belief that, given the stochastic specifications for both assets, it is possible that the payoffs be typically on the positive side. Nevertheless, the results are extremely sensitive to the parameters we chose and exhibit a very high variability from one scenario to the other. Thus, it would be rather dangerous to advance a general opinion based only on these simulations.

The conclusion, however, is that stochastic volatility is a major issue and should not be ignored in the pricing of exotic options.

Price bias of an spread option as a percentage of the Black-Scholes price for varying values of initial price of asset 1, asset 2, volatility of volatility of asset 1 and volatility of volatility of asset 2. Parameters of the option: a=1, b=-1, σ_1=15% , σ_2=15% , K=0, ρ=0.82, σ_1*=15% , σ_2*=15% r=0, T=1/2. a=10 (speed of mean reverting process)

Vol of vol asset1	Vol of vol asset1	Initial price asset1	Initial price asset2				
1	1		0.9	0.95	1	1.05	1.1
		0.9	8.16%	23.08%	76.25%	120.81%	2487.48%
		0.95	2.71%	6.46%	21.12%	37.69%	154.33%
		1	0.98%	2.36%	7.39%	24.82%	45.71%
		1.05	0.36%	0.44%	3.53%	6.00%	26.46%
		1.1	0.10%	0.47%	0.98%	3.72%	9.63%
2	2		0.9	0.95	1	1.05	1.1
		0.9	14.61%	45.62%	174.12%	1727.32%	4083.56%
		0.95	5.59%	20.18%	75.43%	156.36%	634.33%
		1	2.61%	7.35%	24.61%	79.76%	141.24%
		1.05	1.03%	3.61%	6.45%	21.46%	52.68%
		1.1	0.31%	1.22%	3.33%	6.94%	17.09%
3	3		0.9	0.95	1	1.05	1.1
		0.9	23.38%	93.07%	271.16%	1238.62%	3847.74%
		0.95	12.05%	25.52%	72.61%	353.91%	1229.00%
		1	5.60%	15.16%	25.49%	75.95%	272.21%
		1.05	1.63%	3.91%	9.33%	27.56%	76.71%
		1.1	0.73%	3.54%	5.28%	11.12%	31.41%

Figure 2.6:

Chapter 3

Hedging: plain vanilla options vs. exotics

3.1 Classical theory on hedging options

3.1.1 Review of the literature

Any investor that sells a contingent claim is faced with two challenging questions: "How should I price the claim" and "How should I deal with the risks incurred by my position"? The previous two chapters have attempted to answer the first question, so we shall now address the second. The answer dates back to the seminal papers by Black-Scholes (1973) and Merton (1973). Under a set of assumptions, they showed that it is possible to replicate the payoff of a derivative security by a dynamic trading strategy so that the risk is eliminated.

Specifically, a dynamic trading strategy consists of units invested in the risk-free asset and units in the underlying asset of the derivative so that the payoff of the option is replicated. The delta here represents the rate of change of the option price with respect to the price of the underlying. By holding units of the asset, the investor offsets the delta of the option position and, thus, holds a "delta neutral" portfolio. Of course, delta changes from one period to the other, so the investor must rebalance his portfolio periodically in order to stay hedged. The alpha will insure that trading the underlying asset is possible through borrowing (or bond trading). This strategy is called "self-financing" if its cumulative cost will be 0, provided that the initial premium of the option is also invested.

Elegant and simple, the hedging strategy above relies on several crucial hypotheses. First, the trader must have the possibility of trading continuously and he must be a "price taker" or "small" compared to the size of the market. The latter condition is subject to some reserves, while the former is totally inapplicable. Second, markets must be complete, with no frictions such as taxes and other transaction costs. In reality, every rebalancing implies a cost in the form of commissions or bid-ask spreads. Thus an increase in the frequency (and efficiency) of hedging must be traded-off against costs that can become prohibitive. Third, the Black-Scholes-Merton model assumes constant volatility, but, as we have already seen, this is not supported by empirical evidence. In general, the implied volatility tends to increase as the stock price goes down and decreases in the other case. It can be shown that with a change in the implied volatility of around 0.5% the hedging performance of Black-Scoles deteriorates sharply. Fourth, there are other aspects that may pose dangers to such a hedging approach: jump movements in the stock price, uncertainty regarding future interest rates or future dividends, limitations on the amount of borrowing and/or on short-selling, indivisibility of securities, etc.

Extensive academic literature documents shows how these issues must be handled. On a general basis, the approaches to defining a hedging strategy could be classified into two categories:

- methods that are preference-free, i.e. they do not depend on expectations and subjective probabilities; they are, as Heinzl (1999) explains it: "time-local in the sense that the hedging strategy depends entirely on the composition of the hedge portfolio and the asset prices at that time, with no regard for future and past";

- methods that attempt to optimize some quantity (maximize expected utility, minimize risk) and that, according to the same author: "depend on the subjective assessment of agents of the security market's future".

The first approach is strongly related to the notions of no-arbitrage and completeness in financial markets. Under these conditions, the value of the replication strategy must be equal to the price of the contingent claim and this strategy will be unique. In an incomplete market or in an imperfect market, it may not be possible to replicate perfectly the derivatives, but only “close enough”.

The problem of hedging options in discrete time has been the subject of several papers, among which: Boyle and Emanuel (1980), Galai (1983), Figlewski (1989), Robins and Schachter (1994), Kabanov and Safarian (1997). Boyle and Emanuel (1980) notice that the probability distribution of hedge returns is affected by discrete rebalancing, so the methodology used in empirical tests of option models should be changed. Galai (1983) finds that the return from discrete adjustment of the hedge does not have a significant effect on the mean return compared to both the return from deviations of the real price from the model price and the riskless return on the investment. However, the variance of the total return increases with the holding period. Robins and Schachter (1994) conclude that discrete delta-based hedges perform well only when the risk measure is market risk or when the options hedged are long-term, in-the-money ones. Otherwise, the hedge ratios must be modified so that the hedge variance be diminished.

An even more important “imperfection” of financial markets is the existence of transaction costs. The total cost of a particular hedging strategy is a function of the frequency of rebalancing. But then rebalancing depends on the particular path followed by the prices, so it is impossible to know from the beginning how large the costs will be. Moreover, the cost structure is not the same across different classes of traders (retail investors, market makers, etc.). In the transaction cost literature, there are two traditional categories of hedging strategies: the ”fixed time strategies” where the hedge portfolio is rebalanced at fixed times and the “tolerance of delta”-based strategies, i.e. adjust the portfolio when the hedge ratio moves outside a predefined interval. The first one has been described by Boyle and Emanuel (1980) or Leland (1985), while the second by Whalley and Wilmott (1993). The latter derive a band around Black-Scholes delta which depends on costs and also on the option gamma.

Leland (1985) finds that the strategy of replicating must depend on the level of transaction costs and on the exact time between portfolio adjustment. The modification of the actual prices due to transaction costs takes the shape of a simple modification in volatility. So, the strategy is, in the end, just a “modified” Black-Scholes. Kabanov and

Safarian (1997) analyze Leland's model and find that constant (proportional) transaction costs can not lead to a small hedging error when the time interval between rebalancing is diminished. This will be the case, however, if the transaction costs are themselves decreasing. Boyle and Vorst (1992) derive self-financing strategies that perfectly replicate a long (short) call, provided discrete rebalancing and transaction costs. The stock price evolution is given by a multi-period binomial model. As a result, they are able to derive bounds on option prices which are larger than those derived by Leland.

Now we shall refer to the case of incomplete markets and we shall shortly mention three approaches to derivatives analysis that are practically based on hedging arguments, namely: super-replication, local risk minimization and expected utility maximization. The first one tries to find "the cheapest self-financing trading strategy that yields a terminal payoff no smaller than the payoff of the derivative one wants to cover" [1]. Classical references on this issue are the papers by: Bensaid, Lesne, Pages and Scheinkman (1992), Edirisinghe, Naik and Uppal (1993) or Mercurio and Vorst (1997). The local risk minimization is a mean-variance criterion for pricing and hedging derivatives which was introduced by Follmer and Sondermann (1986) and thereafter developed by Follmer and Schweizer (1991) and Schweizer (1995). The idea is to determine a trading strategy in the underlying asset that reduces the risk of the derivative to some "intrinsic component". It means that some part of the risk can still be hedged in a classical manner, but for the remaining part, economic equilibrium arguments or concepts from insurance pricing are necessary. The last approach, maximization of utility, combines transaction costs and utility theory. The seminal paper in this field belongs to Hodges and Neuberger (1992) while Davis, Panas and Zariphopoulou (1993) make additional improvements to the underlying theory. The idea is that the investor knows that the hedging is costly and therefore must define ex-ante some utility function that characterizes his risk preferences. Usually, the chosen function in these models is the exponential utility one because it has the nice advantage of constant risk aversion. The result is a no-transaction region in which the investor must remain. Below or above this region, transactions must be done in order to reach its boundaries.

[1] Frey, R., (1997) - "Derivative Asset Analysis in Models with Level-Dependent and Stochastic Volatility", CWI Quarterly, 10, 1-34

The natural question would be: which strategy is the best? Of course, there is no simple answer since different people will have completely different views on what "best" means. Any search for a solution to this problem must start with a definition of the optimality criteria.

3.1.2 The base case: a European call option

Our research is oriented on the impact of different "imperfections" on the quality of the hedge as shown by means of numerical simulations. As such, we are in the same line of thought with authors like: Figlewski (1989) or Mohamed (1994).

Figlewski (1989) uses Monte Carlo to determine how discrete rebalancing, indivisibilities, transaction costs and uncertain volatility may affect the traditional delta-hedge of a call option. Discrete rebalancing leads to a 0 mean return from hedging, as expected, but the standard deviation is different from 0. Then, the hedge is implemented considering that the "true" volatility of the stock is different that the one implied by the market. A trader that infers the correct volatility can take advantage of the mispricing through a dynamic hedging strategy. The reverse, a mistake in estimating volatility, is dangerous since it affects the hedge ratio at each rebalancing. This error is particularly relevant for out-of-the-money options which are very sensitive to volatility.

An interesting result is that: "hedging with too high a volatility estimate does not seem to increase risk much at all. However, underestimating the volatility leads to a considerably larger standard deviation. These results suggest that, in trying to cope with uncertainty about volatility, it might be appropriate to compute the hedge ratio for out-of-the-money options using a higher volatility than what the trader expects in the future, on the grounds that it is less costly to err on the side of overestimating than underestimating volatility for these options"[2]. Furthermore, the transaction costs are assumed different for retail investors and market makers and this will lead to extreme differences in the efficiency of hedging.

Mohamed (1994) compares the performance of four hedging strategies: Black-Scholes delta hedging at fixed intervals, Leland's delta hedging at regular intervals, hedging using

[2]Figlewski, S. (1989) - "Options Arbitrage in Imperfect Markets", Journal of Finance 44, 1289-1311

a fixed band around Black-Scholes delta and, finally, the Whalley-Wilmott band around the Black-Scholes delta. Monte Carlo simulations are performed; the hedging errors and transaction costs are traded-off against each other. Leland's strategy seems to outperform the simple Black-Scholes one, as expected. Both fixed revision period strategies are better than the hedging with a pre-established band around delta. However, the best one seems to be the Whalley-Wilmott one based on utility maximization.

We have chosen the case of a simple European call option in order to outline the details of the simulation process and to have a benchmark for comparison of hedging performance. In other words, we shall compare the dynamic hedging of the call with the dynamic hedging of an exotic option and we shall see if the latter performs better or worse. Obviously, if hedging the exotic option is at least as efficient as hedging a simple call, any financial institution should implement this strategy. We are aware that there are other possibilities to assess the hedging performance, for example, a value at risk approach or the 95% worst case loss from all the simulated hedging errors[3]. However, we consider that since, in practice, the hedging errors for a call are generally accepted, then this should be the case for exotic options that exhibit the same (or fewer) errors.

Our approach will be also based on Monte Carlo and will account for discrete-time adjustments, transaction costs (fixed and proportional) and stochastic volatility. As a result, we have analyzed for possible cases:

- discrete rebalancing, no transaction costs and constant volatility;
- discrete rebalancing, with transaction costs and constant volatility;
- discrete rebalancing, no transaction costs and stochastic volatility;
- discrete rebalancing, with transaction costs and stochastic volatility.

The implementation is done in Gauss. For the first case, "Discrete Rebalancing, No Transaction Costs and Constant Volatility", we have simulated 10000 underlying stock prices paths over 180 days. The stock did not pay any dividends and, except for

[3]Mohamed (1994) has sorted the 1000 simulated losses in ascending order and has chosen the 950th entry as it indicates: "the loss which would only be exceeded with 5% probability"

the discrete rebalancing, the market was assumed "perfect". The price process was the traditional geometric Brownian motion:

$$dS = \mu S\, dt + \sigma S\, dw \tag{3.1}$$

Accordingly, the price series was obtained by applying Ito's Lemma to the natural log of S, namely:

$$S_t = S_{t-\Delta} e^{\left(\mu - \frac{1}{2}\sigma^2\right)\Delta + \sigma\epsilon\sqrt{\Delta}} \tag{3.2}$$

where

$\epsilon \sim N(0,1)$random

S: stock price

μ: drift of the price process

σ: volatility of prices

Δ: time interval between observations (here 1 day)

We have assumed that the drift of the price process is the risk-free rate. Of course, this is not valid if the drift rate of the stock cannot be hedged. But, following the approach of Mohamed (1994), we consider, just like him, that: "few option writers would hazard a guess for the drift rate for the life of the option. Is there a discrete time strategy that immunizes the writer from variations in μ?"

We considered the case of an option writer that hedges his position in the classical manner. S/he sells an European call option with time to maturity of 180 days and immediately purchases delta amount of the underlying stock. Delta is, of course, the first derivative of the option's price with respect to the underlying stock and it is given by the formula:

$$\delta = N(d_1) \quad \text{where} \quad d_1 = \frac{\ln\left(\frac{S}{k}\right) + \left(r + \frac{\sigma^2}{2}\right) T}{\sigma\sqrt{T}} \tag{3.3}$$

Figure 3.1. represents a possible path for the option price and its replicating portfolio under the previous conditions.

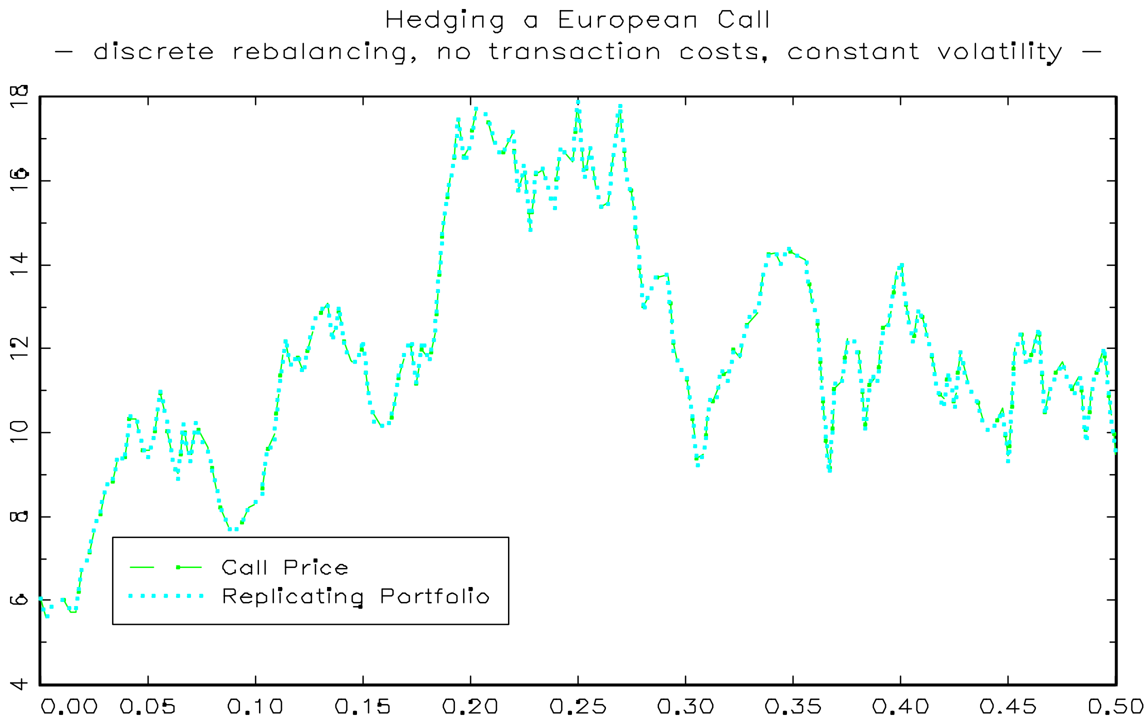

Figure 3.1:

The purchase is partially financed by the option's premium, computed according to Black-Scholes formula, and the rest is borrowed at the risk free rate. The parameters used for this simulation were:

- starting value of the stock price: $S_0 = 100\$$;
- strike price: $K = 100\$$;
- time to maturity: $T = 180$ days;
- volatility of underlying asset: $\sigma = 20\%$;
- risk free interest rate $r = 7\%$;
- number of observations: $N = 180$, i.e. daily observations;

The dynamic trading strategy for each price path consisted in recalculating the delta in each subsequent day and in rebalancing the amount of stock accordingly. Also, at each moment in time, both the value of the hedging portfolio and the option price were computed.

Under these almost perfect market conditions, the discounted cost of hedging the option should be equal to its theoretical value. In our case, we have taken, for each simulation, the difference between the last value of the hedge portfolio ($V(T)$) and the last value of the call ($\text{Call}(T)$), then we have computed the mean of these values. Thus, we have obtained the average hedging cost (HedgeCost) and finally we have divided this value by the initial call value ($\text{Call}(t)$) in order to determine the percentage hedging cost.

$$\text{HedgeCost} = \frac{V(T) - \text{Call}(T)}{\text{Call}(t)} \tag{3.4}$$

Furthermore, we have assessed the variability of the hedge ($\sigma_{\text{HedgeCost}}$) by computing the standard deviation of the cost of hedging. Due to the division by the same initial value of the call, this variability is expressed in percentage terms.

$$\text{HedgeVar} = \sigma_{\text{HedgeCost}} \tag{3.5}$$

For the above parameters and 10000 simulations we have obtained an average hedging cost of -0.0378% and a variability of the hedge of 4.9524%. The 95% confidence interval for the hedging cost was: $(-0.0438\%, 0.1195\%)$. As expected, the average hedging cost is not significantly different from 0.

Next we tried to assess the effect of rebalancing at different time intervals on both the average hedging cost and its standard deviation. So, we have supposed that delta is recalculated and the strategy is implemented every day, than every two days, etc. The results for the average hedging costs and its confidence interval are in figure 3.2.

We have repeated the procedure to see how the hedging cost variability is affected and we have obtained the following results (see figure 3.3).

The conclusion is that if rebalancing takes place less often, the average hedging cost will increase and its variability also.

We must be aware that the strategy presented above is totally unrealistic since it only accounts for one single "imperfection" of the financial market: the impossibility of continuous rebalancing. It is the reason why we are developing the second case, "Discrete Rebalancing, with Transaction Costs and Constant Volatility". We present the effects of incurring transaction costs each time we rebalance our portfolio. Specifically, the costs

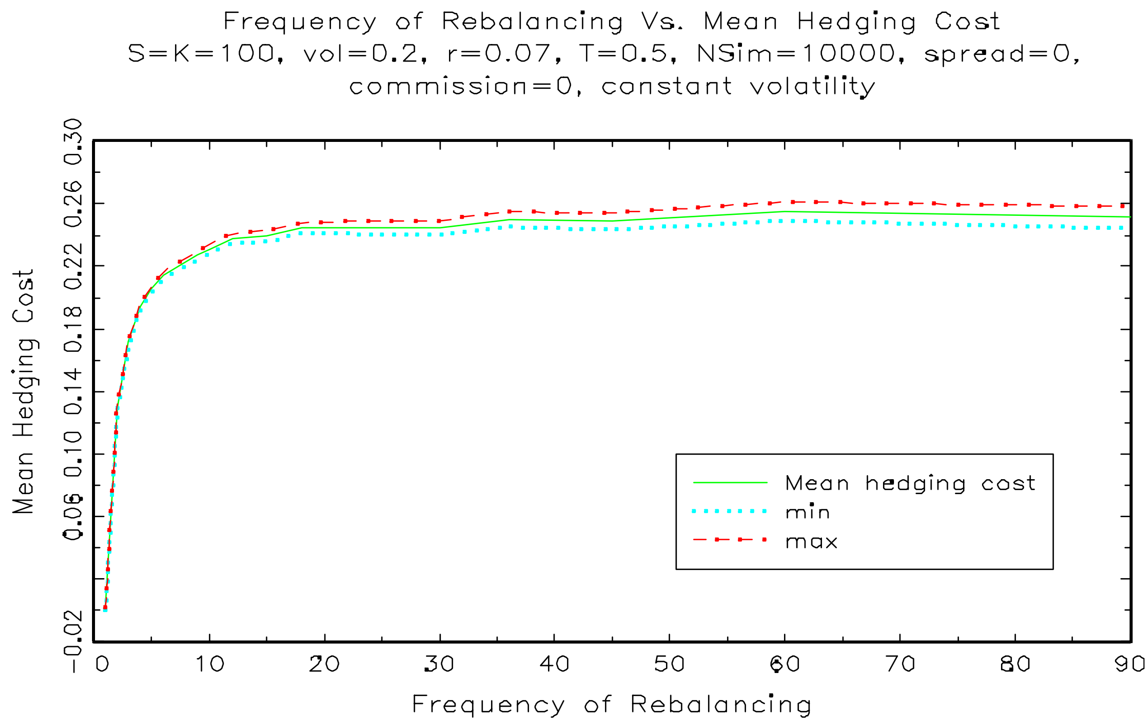

Figure 3.2:

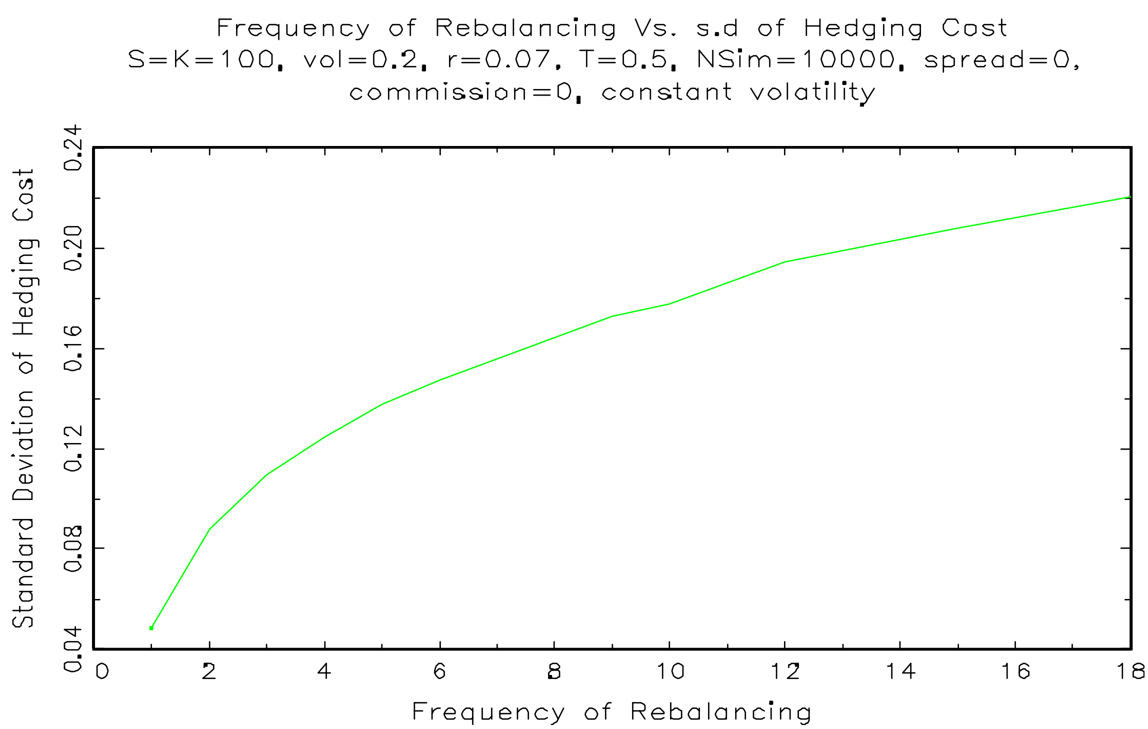

Figure 3.3:

are a fixed bid-ask spread of $\frac{1}{4}$ and a commission on the total value of the transaction of 0.05%. The parameters and the simulation of prices of the previous model are maintained, but the computation of the value of hedging portfolio is modified. At each time step, the delta is computed with the help of the simulated price, like before. But when there is a need to buy additional stock, the price paid is higher than the simulated one with $\frac{1}{8}$ and when stock must be sold, the price received is smaller than the simulated one with the same $\frac{1}{8}$.

Performing 10000 of simulations, with rebalancing each day and transaction costs, we have obtained an average hedging cost of -12.5204% and a corresponding variability of hedge of 6.5073%. The 95% confidence interval for this hedging cost was: $(-12.6277, -12.4130\%)$. It is easy to see that the hedging cost has risen to approximately 12% of the initial call value, which is understandable given the existing transaction costs. However, the variability of the hedge is only slightly modified compared to the case with no transaction costs.

As shown by Figlewski (1989), the transaction costs are not the same for a market maker and for a retail investor and this can lead to very different results in the hedging performance. For comparison, we have studied the effect of changes in the commission on the average hedging cost. The results are represented graphically (see figure 3.4).

The conclusion is that transaction costs lead to a significant underperformance of the hedging strategy, even for a market maker. The retail hedger may find this strategy completely prohibitive since his costs and their standard deviation will be very high. A decrease in the hedging frequency can be dangerous for any of them even though it might be tempting due to the decrease in costs.

In order to assess the effect of changes in commission on the variability of the hedge, we have looked at the standard deviation of the hedging costs for different levels of commission (see figure 3.5).

It is easy to notice that variability of the costs is only slightly affected by the increase in costs which is normal.

The third case, "Discrete Rebalancing, No Transaction Costs and Stochastic Volatility" represents an attempt to isolate the effect of stochastic volatility provided that there

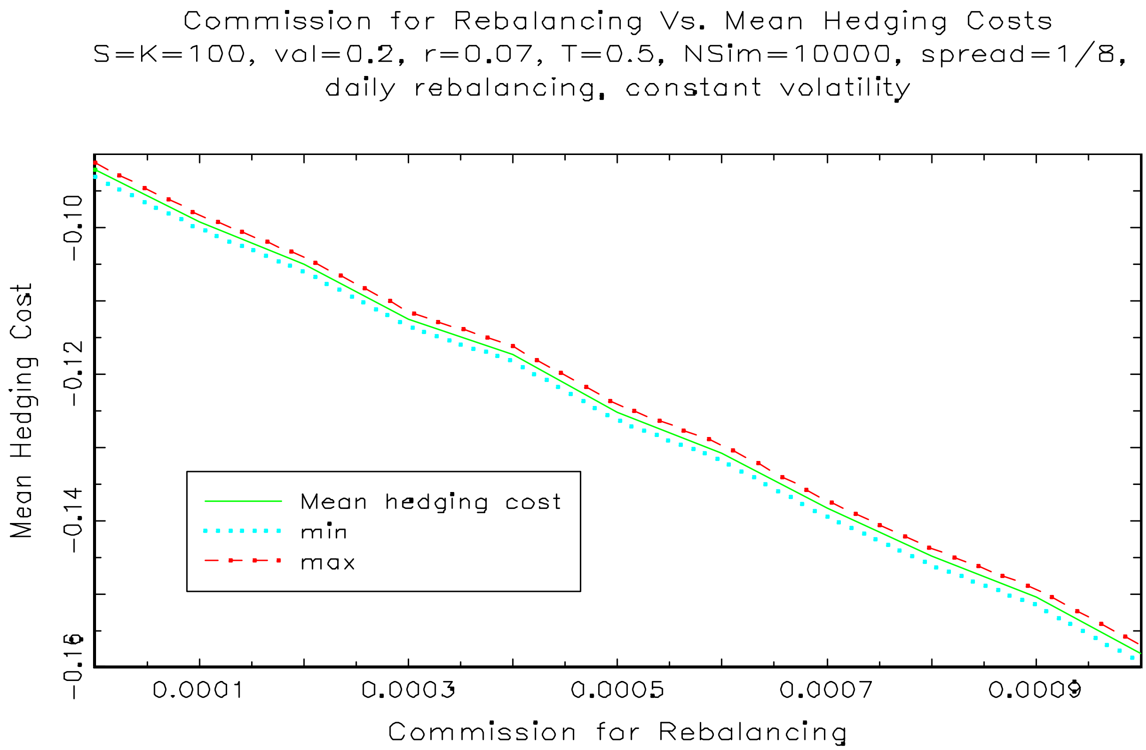

Figure 3.4:

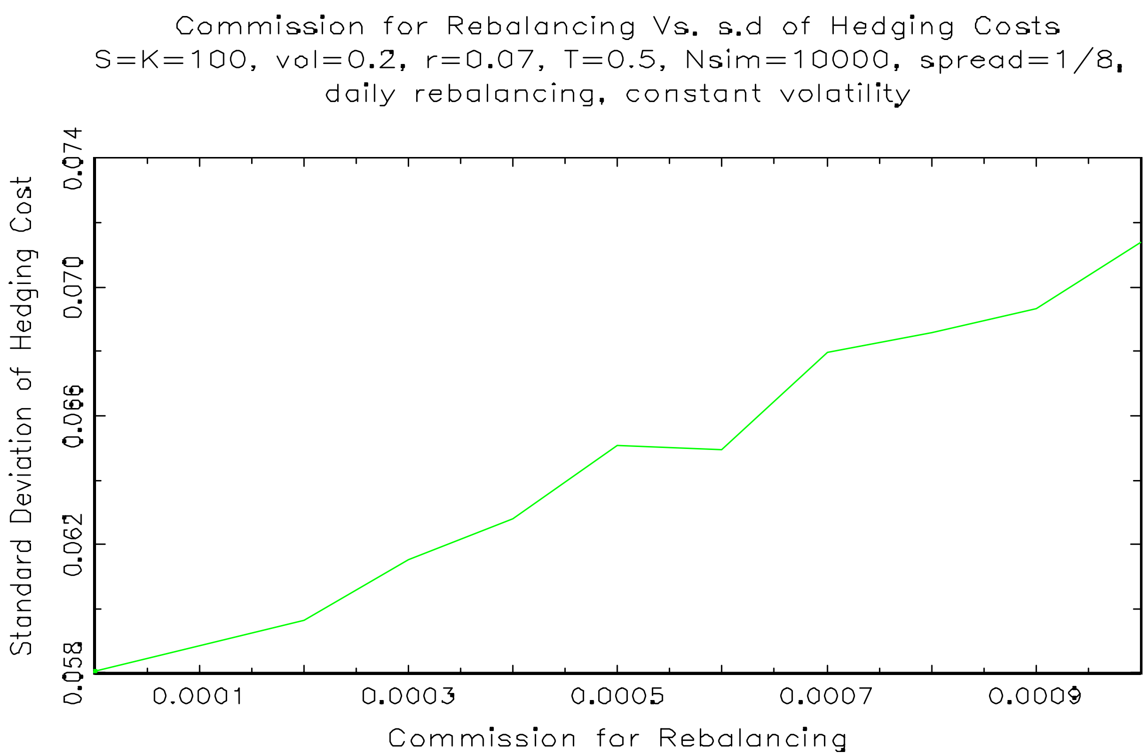

Figure 3.5:

are no costs. However, in order to stay realistic, the rebalancing takes place discretely. The stochastic process followed by variance is the previously presented mean-reversion process, that is:

$$dV = a\,(\bar{\sigma} - \sigma)\,V\,dt + \xi V\,dz \tag{3.6}$$

When this process is considered, we assume first that there is 0 correlation between the underlying asset price and volatility. Then, the parameters are fixed similar to those chosen by Tompkins (2002), namely: the rate of mean reversion $a = 16$, the volatility of volatility $\xi = 1$ and the long-term mean of volatility was set to 20% annually. Using these fixed values, 1000 paths for the volatility were generated with the help of random standard normal variables. These volatilities were used to obtain 1000 price paths. For each such path, the dynamic hedging was implemented as usual. Delta was calculated at each step according to Black-Scholes and having as an input the corresponding simulated volatility. The theoretical call prices at each rebalancing were also obtained with the help of the stochastic volatilities.

The average hedging cost was 1.6932% and the standard deviation 12.7970%. The interval of confidence was: $(1.4820\%, 1.9043\%)$. Even though including the stochastic volatility had no major effect on the hedging cost, the variability of the hedge has significantly increased. Hull & White (1987) demonstrated that the theoretical value of a call is nothing else than the Black & Scholes formula with the constant volatility replaced by the average realized volatility over the life of the option. In our case, since the expected average volatility will be approximately 20%, i.e. equal to the constant one, we can easily understand why the expectation of the hedging cost must be close to 0.

The results are extremely sensitive to the parameters chosen. For example, we could move the volatility of volatility and see how the hedge performs. Hull & White (1987) have performed several empirical studies according to which this parameter can only takes values between 1 and 4. We have chosen the values of 1, 2 and 3 for the volatility of volatility and combined this change with rebalancing at different time intervals. The graph 3.6 illustrates the fluctuations in hedging cost when the volatility of volatility changes.

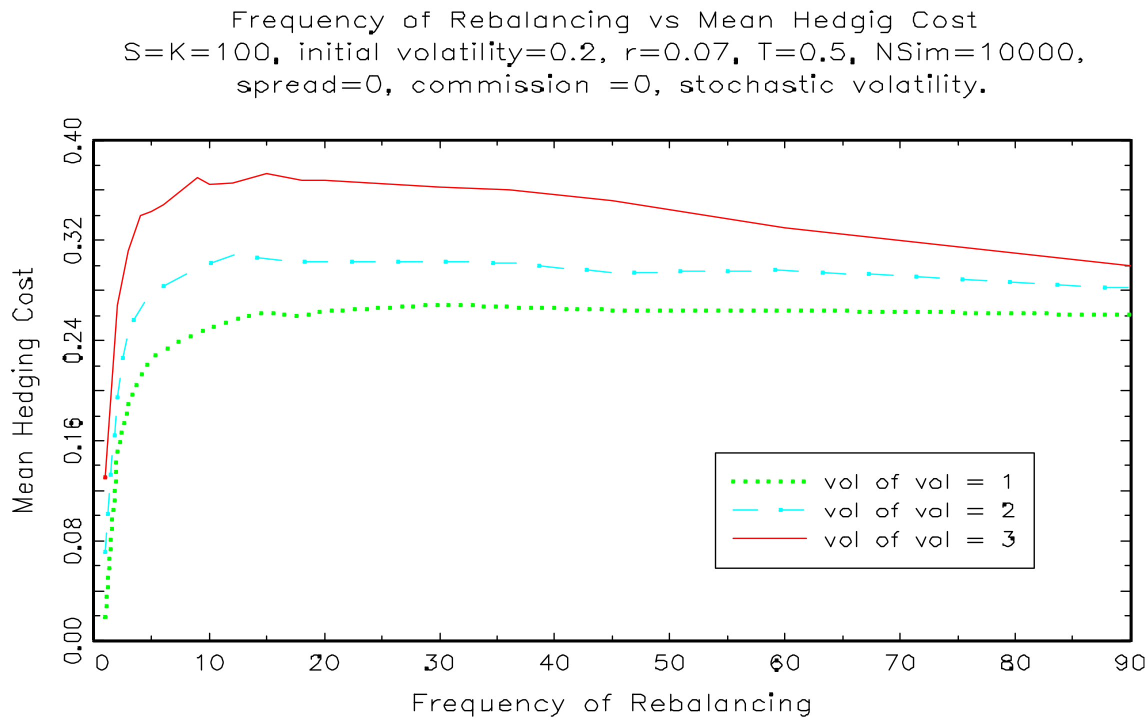

Figure 3.6:

It becomes obvious that the higher the volatility of volatility, the worse the performance of the hedge. The effect on the standard deviation of the hedge is similar, as illustrated in figure 3.7.

One last case will account for all the major "imperfections" of the financial market: "Discrete Rebalancing, with Transaction Costs and Stochastic Volatility". The application will involve the simulation of 10000 volatility paths and, with their help, of 10000 price paths Rebalancing will be done as usual, but the prices will be higher or smaller for buying and selling, respectively. The commission of 0.05% will have to be paid on the total amount of transaction at each rebalancing. Finally, the results obtained were: the average hedging cost was -11.2284% and the corresponding standard deviation was 13.8052%. The interval of confidence for the hedging cost was $(-11.9488\%, -10.5081\%)$. The call price evolution and its replicating portfolio at each step are represented in Figure 3.8. The negative effects of market imperfections become obvious.

The inclusion of transaction costs generates a significant increase of both the average hedging cost and of its variability. Actually, the impact is more visible on the side of the average hedging cost, while the standard deviation is similar to the one in the previous

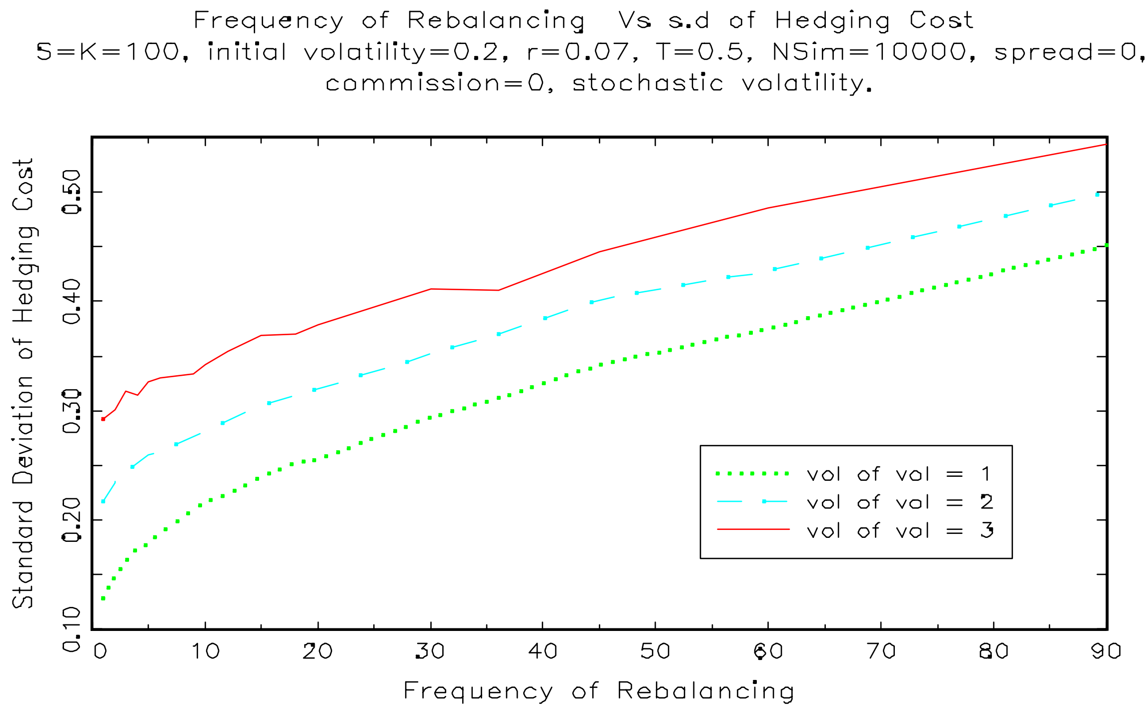

Figure 3.7:

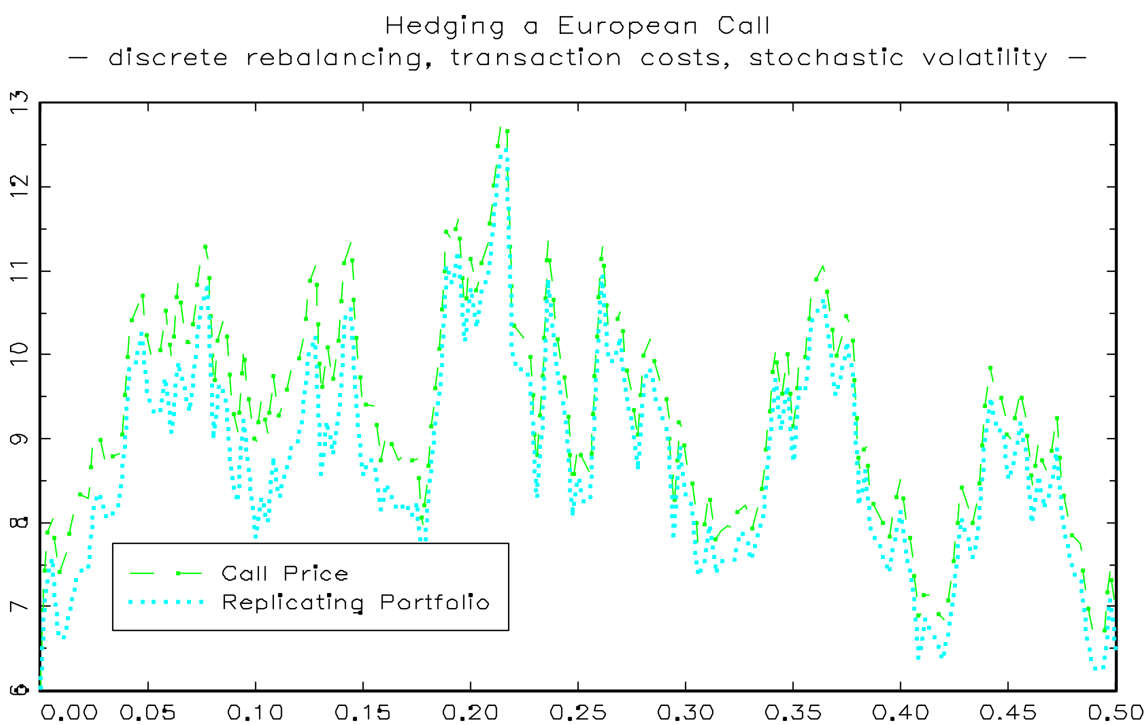

Figure 3.8:

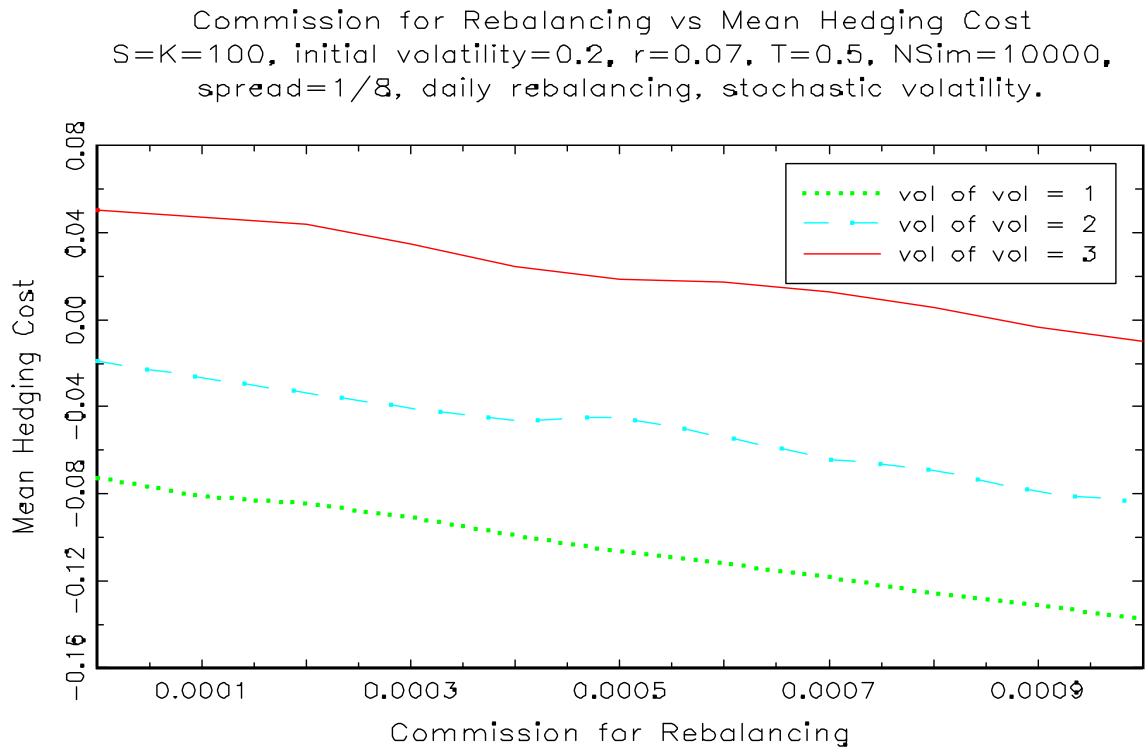

Figure 3.9:

case (stochastic volatility, but no transaction costs). We have tried to understand the effect of changes in the parameters of the volatility process combined with different transaction costs. For this reason, we have chosen the same three values for the volatility of volatility, as before, and gradually increased the commission.

The hedge performs worst for very high transaction costs , while the effect of volatility of volatility in this case is rather ambiguous (see Figure 3.9). As already shown, the standard deviation of the hedge is only slightly affected by changes in commissions, but it is affected by the volatility of volatility. The higher this parameter, the bigger the standard deviation and the more dangerous this hedging strategy (see figure 3.10).

The procedure described above can be implemented for any exotic option as long as the delta (or the deltas in the case of correlation options) are known. However, it may be the case that dynamic hedging is not the best solution and static hedging will be used instead. Nevertheless, the simple case of a European call represents a benchmark of comparison and it helps to understand the foundations of our simulation approach.

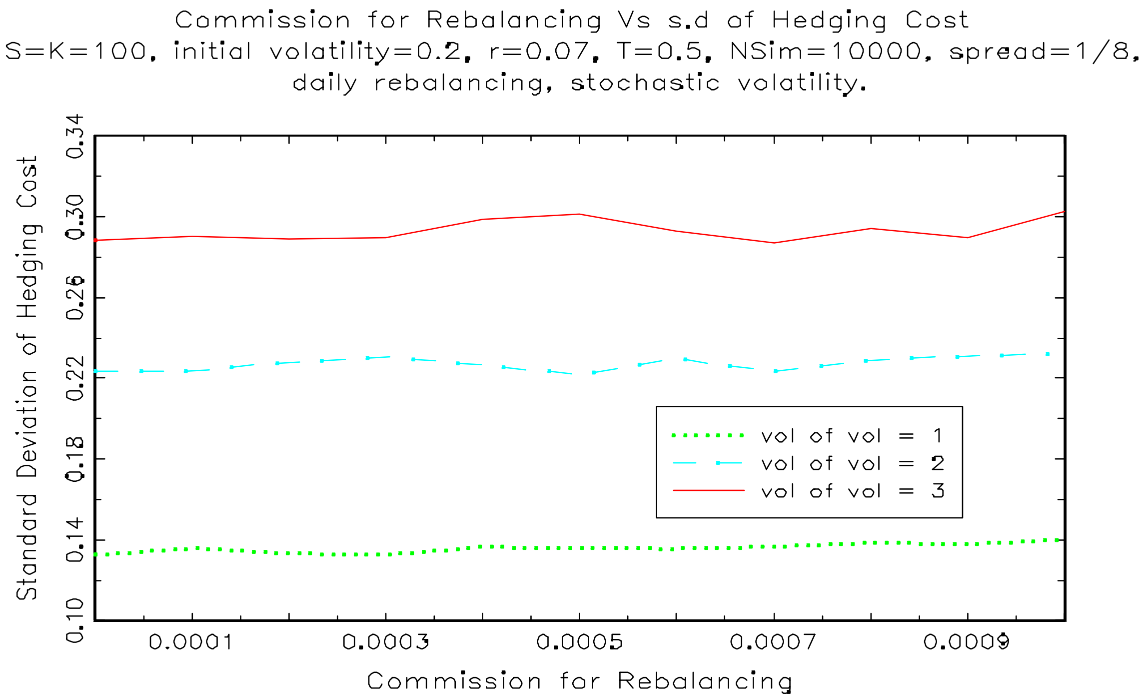

Figure 3.10:

3.2 Hedging strategies for exotic options

In a "perfect market", pricing exotic options is possible, as we have seen, but hedging them is a completely different issue. There are two general approaches to hedging exotic options: static hedging and dynamic hedging. It can not be said which one is dominant on a general basis because their efficiency is different for different types of options as we shall see.

Exotic options can be hedged "statically" using a portfolio of standard options or other products which remains unchanged until the expiration of the exotic claim. Such a strategy avoids transaction costs from rehedging and this is why it may be preferred especially in illiquid markets: "market makers generally try to hedge exotic option positions (arising from customer deals) using plain vanilla options which can be of different times to expiry than the exotic options. In this way, gamma and vega risk exposure are passed from the very illiquid exotic options book to the more liquid plain vanilla options book. [...] Remaining risk exposure, which cannot be hedged in this way, is kept on the exotic options book by holding the contract to maturity"[4]. The main disadvantages of

[4]Moessner, R. (2001) - "Over the Counter Interest Rate Options", Research Paper 1/2001

this approach are that it requires a relatively high amount of income from the trade to be spent paying bid/offer spreads and markets are often illiquid.

Dynamic hedging has been detailed above for a simple call option and it is similar for exotic options. Practically, the option is hedged against directional price movements of the underlying asset(s), by buying and selling delta(s) of this/these underlying(s). This is, by definition, a so-called "delta hedging" and it will be the type of hedging that we shall apply to an exchange option in the next chapter. However, the dynamic hedging is a wider term referring to: "sticking to a minimum Greek exposure and rebalancing continuously to achieve a certain neutrality. [...] Dynamic hedging concerns all the Greeks in the book. It starts with the rebalancing of the deltas (as the market moves or as the delta bleeds with time). As Gammas change, it involves the adjustment through options to reduce or increase gammas and the consequent time decay. As markets move, rho needs to be adjusted and so on."[5]

Next we shall detail the dynamic hedging approach for an exchange option by means of simulations in Gauss.

3.2.1 Hedging an exchange option under transaction costs and stochastic volatility

We have chosen to detail a dynamic trading strategy for exchange options since they are the basic type of correlation options. Actually, many other correlation options can be transformed into exchange options and analyzed as such. Furthermore, they are among the most complicated exotic options since their value depends on two underlying assets, not just one.

As far as the hedging of these products is concerned, Margrabe (1978) proposed a solution that involves the deltas of both assets. Maintaining a position in both underlying assets implies that correlation between them is of major interest. It is the reason why these options have been called "first order correlation dependent options". Moreover, since correlation is not a traded asset, it is extremely hard, if not impossible, to think of a static hedge.

[5] Nassim Taleb (1997) - Dynamic Hedging, Managing Vanilla and Exotic Options, Wiley&Sons

The results of the dynamic hedging procedure for a European call will serve as benchmark. We consider an exchange option on two stocks, S_1 and S_2 that has been sold. The first step is to establish the parameters of the price processes S_1 and S_2, i.e.:

- starting values of the stocks prices: $S_{01} = S_{02} = 100\$$;
- time to maturity: $T = 180$ days;
- volatilities of underlying assets: $\sigma_{01} = 20\%$ and $\sigma_{02} = 15\%$;
- risk free interest rate $r = 7\%$;
- no dividends: $g_1 = g_2 = 0$;
- number of observations: $N = 180$, i.e. daily observations;

We assumed that there was a correlation of 50% between the price processes, but no correlation between the prices and their corresponding volatilities. To simulate two dependent processes, we can use several possibilities. First, the stochastic process for the second asset can be obtained from the Wiener process of the first asset, by the formula:

$$w_{t_2} = \rho w_{t_1} + \sqrt{1-\rho^2} w_t \tag{3.7}$$

where

w_t: random draw from a $N(0,1)$

w_{t_1}: random draw for the Wiener process of the first asset

ρ: correlation between price processes

Second, it is possible to perform a Cholesky decomposition of the correlation matrix and use it to generate correlated random numbers. This last possibility is more general and it is easy to implement in Gauss since there is are specific commands for such a purpose. Moreover, Cholesky decomposition can be very useful for exotic options based on multiple assets such as basket options or complex rainbow options.

The dynamic hedging of the exchange option is different than the one for a simple call, since we have two assets involved, and therefore two deltas. The first delta is the derivative of the exchange option price with respect to the first asset. In order to obtain it, we have used the pricing formula from Zhang (1998), thus:

$$\frac{\partial \text{PriceExch}}{\partial S_1} = e^{-g_1 T} N(d_1) \tag{3.8}$$

where

$$d_1 = \frac{\ln\left(\frac{S_{t_1}}{S_{t_2}}\right) + \left(g_2 - g_1 + \frac{1}{2}\sigma a^2\right) T}{\sigma_a \sqrt{T}}$$

$$\sigma_a = \sqrt{\sigma_{01}^2 - 2\rho\sigma_{01}\sigma_{02} + \sigma_{02}^2}$$

The second delta gives the sensitivity of the exchange option price to the second asset, it is always negative and we have calculated it in the same manner to obtain:

$$\frac{\partial \text{PriceExch}}{\partial S_2} = -e^{-g_2 T} N(-d_2) \tag{3.9}$$

where

$$d_2 = d_1 - \sigma_a \sqrt{T}$$

The hedging procedure started, as for a call, with buying a certain amount of the first asset and selling another amount of the second asset according to their specific initial deltas. The exchange option premium was determined with the Margrabe (1978) formulas and was considered as part of the initial investment. The excess cash needed or held after these operations was borrowed or lent, respectively, at the risk free rate. Each subsequent day, the portfolio was rebalanced according to the new deltas and the values of the exchange option and of the replicating portfolio were recomputed for comparison. The formulas for the hedging cost, its interval of confidence and the standard deviation of the cost were similar to the ones used for the call.

In order to be consistent with the approach used for a simple call, we shall study the same four cases:

- discrete rebalancing, no transaction costs and constant volatility;
- discrete rebalancing, with transaction costs and constant volatility;

- discrete rebalancing, no transaction costs and stochastic volatility;
- discrete rebalancing, with transaction costs and stochastic volatility.

All cases involved the simulation of 1000 price paths for both assets, using the geometric Brownian motion and the parameters stated before. The first case, "Discrete Rebalancing, No Transaction Costs and Constant Volatility" was based on a daily revision of the hedge portfolio. Practically, the two deltas were compared with their previous values and, base on this, a decision was taken regarding buying or selling a specific amount of them. There were four possible combinations of transactions to be performed for rebalancing: (buy S_1, buy S_2), (buy S_1, sell S_2), (sell S_1, buy S_2) and (sell S_1, sell S_2) according to the movement in their deltas.

The average hedging cost obtained after 1000 simulations was 0.1823% and the 95% interval of confidence was: $(-0.1654\%, 0.5301\%)$. The variability of the hedge was 6.6665%. The results are similar to those obtained for a call and this should not be a surprise given the "almost perfect" market conditions that we assumed. It would be interesting to see the effect of discrete rebalancing on the hedge performance, thus we have assumed that the adjustments take place every 2, 3, 4 days, etc. and obtained the Figure 3.11.

The second case, "Discrete Rebalancing, with Transaction Costs and Constant Volatility", outlines significant differences between the base case (the European call) and the exchange option. First, the transaction costs are assumed equal for the two assets: a fixed bid-ask spread of $\frac{1}{4}$. The hedging proceeds like before, except that purchases are performed at the simulated price plus $\frac{1}{8}$ and sales at the simulated price minus $\frac{1}{8}$. Then, for each transaction a commission of 0.05% is paid indifferent of the asset and of the type of trade (buying or selling). The hedging performance is worse than for a call as Figure 3.12 shows it.

The results are: an average hedging cost of -33.1959% with a 95% confidence interval of: $(-33.9104\%, -32.4814\%)$ and a variability of hedge of 13.6937%.

Therefore, we notice that the exchange option is far more sensitive to the transaction costs than a simple call. This is intuitive since we have to trade two assets and incur costs for both. The variability of the hedge is also higher than the one for a European

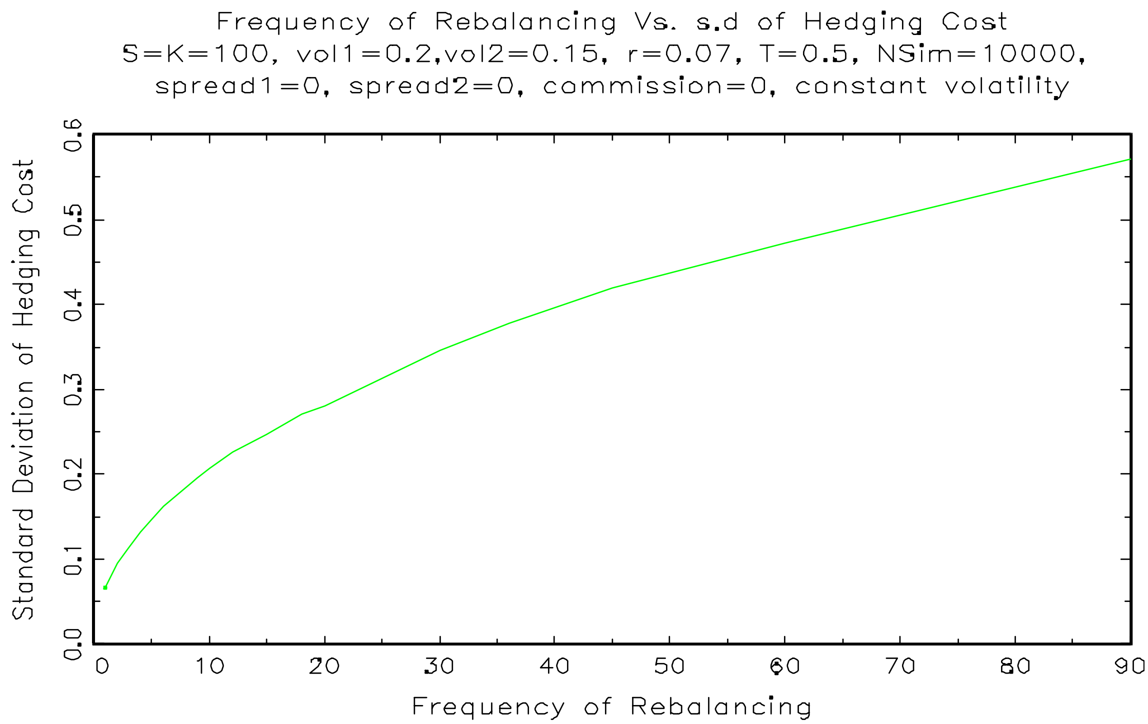

Figure 3.11:

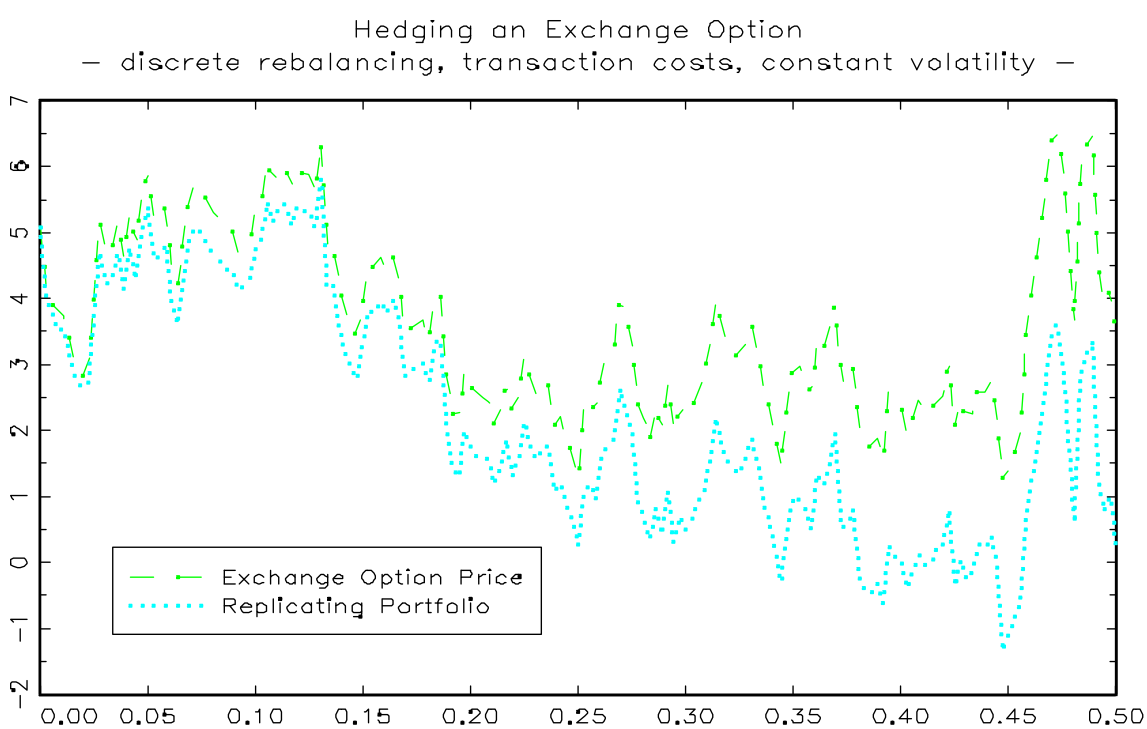

Figure 3.12:

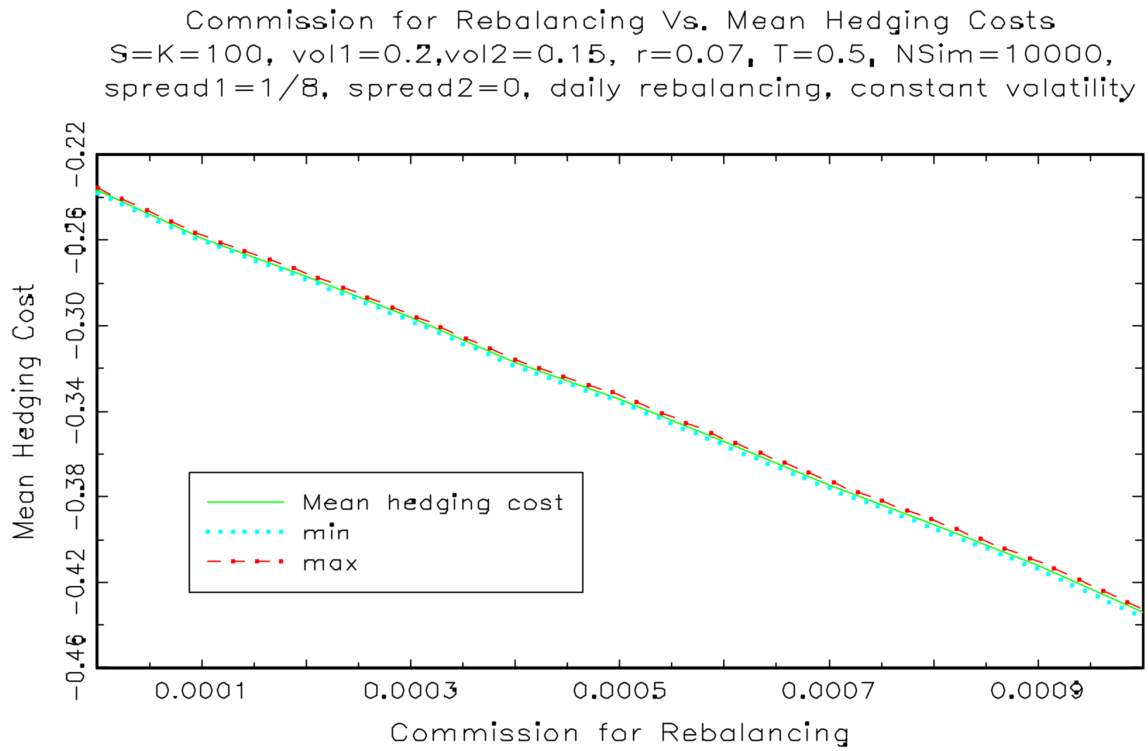

Figure 3.13:

call.

The sensitivity of these results to changes in transaction commissions is illustrated graphically below and it becomes obvious that transaction costs are a major issue when delta hedging such products, as Figures 3.13 and 3.14 show it.

The third case, "Discrete Rebalancing, No Transaction Costs and Stochastic Volatility" complicates the simulations since volatility paths are needed for both assets. So, we assume that volatility follows the Hull & White mean-reverting process for both assets, but we chose slightly different parameters for the processes, namely: the rates of mean reversion are $a_1 = 16$ and $a_2 = 12$, the volatilities of volatility are $\xi_1 = \xi_2 = 1$ and the long-term means of volatility $\sigma_1 = 20\%$ and $\sigma_2 = 15\%$. These parameters will be used for generating 1000 simulations of volatility paths for each asset, then the values at each step will serve as inputs in the price paths. However, the simulation of price paths will involve antithetic variables, so that, finally, there will be 2000 possible paths for each asset. The calculation of deltas and exchange option prices will be performed with the simulated values for volatility, at each time step.

Overall, the results are: an average hedging cost of 0.1030% and with an interval of

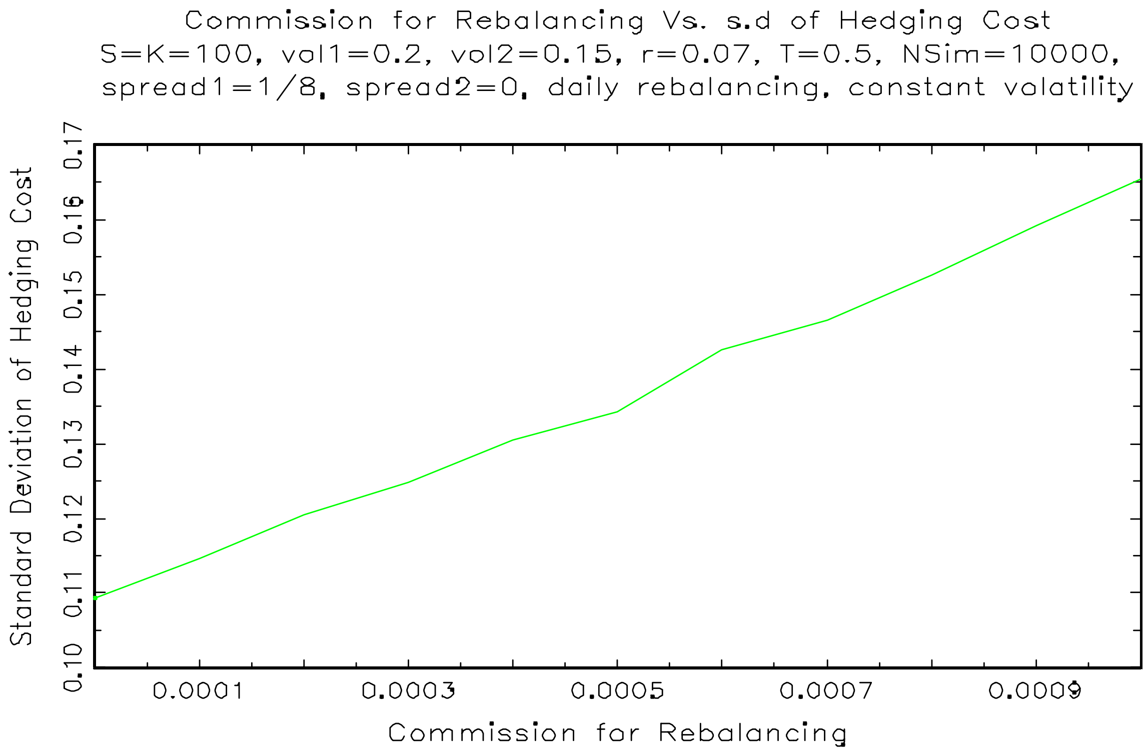

Figure 3.14:

confidence of $(-0.4577\%, 0.6638\%)$, then a standard deviation of hedge of 15.1996%. The hedge variability is therefore similar to the no transaction costs, stochastic volatility case for an European call and, of course, it is higher than for the case of constant volatility.

The last case, "Discrete Rebalancing, with Transaction Costs and Stochastic Volatility", will involve the same 2000 simulation for each asset's price. Following the comparison of deltas with their previous values and the trade direction, costs must be supported. The values of spread and commission will be the same for both assets, i.e. 1/8 and 0.05% and the simulated volatility values will also be used at each step. The results are average hedging cost of -33.0606% and its confidence interval is $(-33.7835\%, -32.3376\%)$, while the variability of the hedge is 19.5954%. We notice an increase in the hedging costs due to the inclusion of transaction expenses. The standard deviation of the hedge is slightly higher compared to the previous case and also somewhat higher than for a call.

The results are sensitive to the volatility of volatility and transaction costs and, in order to understand the impact of this parameter for both volatility processes, we have calculated the mean cost and variability for different values of ξ. (see Figures 3.15 and 3.16)

Mean hedging cost

constant volatility			
spread1=0 spread2=0 commission 0	spread1=spread2=1/8 commission 0.0002	commission 0.0005	commission 0.0008
0.000235515	-0.27665592	-0.3339855	-0.3928001

stochastic volatility					
vol of vol asset 1	vol of vol asset 2	spread1=0 spread2=0 commission 0	spread1=spread2=1/8 commission 0.0002	commission 0.0005	commission 0.0008
1	1	0.004404359	-0.27486062	-0.3346894	-0.3914433
1	2	-0.014111234	-0.28787096	-0.3449162	-0.4015859
1	3	-0.031033175	-0.30689609	-0.3640838	-0.418021
2	1	0.025772924	-0.24640581	-0.3052581	-0.3579204
2	2	0.020766936	-0.25339239	-0.3137678	-0.3658538
2	3	0.011571068	-0.25947808	-0.3114157	-0.3737757
3	1	0.048602463	-0.21053513	-0.2713673	-0.3275675
3	2	0.056436558	-0.20118958	-0.2645644	-0.3231189
3	3	0.069684805	-0.19513374	-0.2566392	-0.3097053

Figure 3.15:

Standard deviation of hedging cost

constant volatility			
spread1=0 spread2=0 commission 0	spread1=spread2=1/8 commission 0.0002	spread1=spread2=1/8 commission 0.0005	spread1=spread2=1/8 commission 0.0008
0.066627061	0.12054742	0.13574363	0.15306008

stochastic volatility					
vol of vol asset 1	vol of vol asset 2	spread1=0 spread2=0 commission 0	spread1=spread2=1/8 commission 0.0002	spread1=spread2=1/8 commission 0.0005	spread1=spread2=1/8 commission 0.0008
1	1	0.15104757	0.18463093	0.19660109	0.20687856
1	2	0.19071036	0.21431622	0.2228428	0.23262827
1	3	0.22825432	0.25406661	0.26576518	0.26616782
2	1	0.24949544	0.2683986	0.27503406	0.27658931
2	2	0.28959198	0.30920105	0.31346799	0.31995442
2	3	0.33578054	0.35500219	0.34957256	0.35762671
3	1	0.3196495	0.32412415	0.32431535	0.34158278
3	2	0.36448162	0.3690503	0.37893511	0.39135285
3	3	0.40831273	0.4199648	0.43391422	0.44143188

Figure 3.16:

The final conclusion is that the dynamic hedging of an exchange option has the same features as the hedging of a call and approximately the same performance. The main difference is that the transaction costs in the case of an exchange option have a stronger effect on hedging than in the case of a call. This effect is understandable due to the trading of two assets, not just one. Overall, we think that if the call is to be taken as a benchmark, there is no evident reason why the dynamic hedging of an exchange option should be rejected.

3.2.2 Static hedging vs. dynamic hedging

As stated in the introduction of this chapter, there are two main approaches to hedging exotic options: static and dynamic. We have illustrated the second one for an exchange option by Monte Carlo simulations. However, there is a significant number of exotic option types for which static hedging is more efficient. For example, path-dependent options often have high gammas, so it may be the case that static hedging is cheaper and less difficult than dynamic hedging. In conclusion, we have attempted to see which approach is best for several exotics and how it can be implemented.

Asian Options

At a closer look, an Asian option is nothing else than a simple call or put, but with a special strike represented by the average (arithmetic or geometric) of observed prices over a given period of time. The computation of this strike is based on closing or settlement prices of the underlying stock and they are cash settled at expiration.

It is a well known fact that dynamic replication of these instruments is easier than for the plain vanilla calls and puts. This is mainly due to the pattern of the delta for these options: although delta tends to be quite volatile at the beginning of the option's lifetime, it will decrease with increasing past averaging period. It simply means that if a large part of the averaging period has already passed, the remaining observations will not have an important impact on the final payoff. As a result the volatility of the price will be reduced, thus also gamma will decrease and delta will be smaller and more stable.

Accordingly, the impact of transaction costs will probably be smaller than for a normal call due to the stability of delta and the less rebalancing required. A similar strategy to the one for an exchange option can be implemented and it will most likely have a similar or superior performance. On the other hand, we must acknowledge the fact that the hedging parameters such as delta will be strongly affected by the observation frequency of the Asian option. For example, Francois (2001) proves that the delta of a fixed strike Asian call is usually overestimated by continuous averaging (and underestimated for puts). This hedging error is usually increasing with time to maturity and decreasing with the option's moneyness. Similarly, gamma for long options tends to be overestimated. Using the continuous averaging option formulas in deriving the Greeks may, thus, be dangerous for hedging.

There is also another possibility of hedging an Asian option, a "semi-static" approach, that consists in buying a simple European option with the same strike, if possible, but with expiration one third of the averaging period. The theoretical values of the Asian and of the standard option will be quite close, so costs will be approximately 0 at the beginning. This strategy may result in some offsetting effects for the gamma and the volatility exposure up to the point where the Asian option becomes easier to hedge. From that point on, however, transaction costs must be supported or, alternatively, the trader can choose not to hedge at all.

Barrier options

Barrier options are "conditional options", namely they depend on whether a certain pre-established has been reached or no by the underlying asset during the option's life. As already shown there are many types of such options, but we shall restrict our hedging analysis to the standard ones: knock-in and knock-out "plain vanilla" options.

It is well-known that these options can be replicated using the traditional delta-hedging approach. The dynamic hedging approach has been implemented for down-and-out call, down-and-in call and up-and-out call by Tompkins (2002). Due to the multiple problems that a dynamic hedging may create for these options, several static hedging approaches have been proposed in papers by: Derman, Kani & Ergener (1994), Carr & Chou (1996) and Carr, Ellis & Gupta (1998). We shall present these alternative hedging strategies for the three options mentioned and, of course, the results can be extended to all other barrier options due to the existing symmetries.

The results obtained when dynamically hedging these three options were very different: the hedging of the down-and-out call exhibited similar performance to a standard call hedging. The reason may be that: "as normal knock-outs have less gamma than European options, there is a strong tendency to rely on delta hedging"[6]. Practically, this barrier option is identical to a standard one and must be hedged accordingly until the barrier is reached; afterwards there is not need for delta-hedging which means that, overall, it may be even easier and less costly to hedge than a normal option.

The down-and-in call may also be dynamically hedged until the barrier is reached, but at that point the delta immediately changes sign and hedging errors may occur. It is the reason why a semi-static hedging strategy has been proposed for this instrument, i.e. delta-hedging until the barrier is breached than purchase of a call option with the same maturity as the barrier option.

The variability of dynamic hedging is the largest for the up-and-out call because this option is in-the-money when it is very close to the barrier, so a large amount of the underlying asset must be bought. However, if the barrier is surpassed this large quantity must be entirely liquidated, and all these trades generate significant hedging costs. So, in general, dynamic hedging is good for "out" options only if they are out-of-the-money just before the barrier is touched. The reverse case, "out" options that are in-the-money close to barrier is dangerous, while for "in" options fully dynamic hedging may raise problems at the barrier point.

Derman, Kani & Ergener (1994) presents a very intuitive method of replicating barrier options, in particular an up-and-out call which is in-the-money close to barrier. So, we have seen before that dynamic strategies may be dangerous for such a product, yet it is possible to almost perfectly replicate it statically using a portfolio of standard call options. This portfolio contains one call option with the same strike price as the up-and-out call which will perfectly imitate the barrier option before being knocked-out. In addition, the portfolio contains many other calls with a strike equal to the barrier level,

[6]Tompkins, R. G (2002) - "Static versus Dynamic Hedging of Exotic Options: An Evaluation of Hedge Performance via Simulation", The Journal of Risk and Finance, 2002

but with different expiration times. The more options in the portfolio, the better the replication. The weights of these options need not be adjusted at any time and they will "replicate the value of the target option for a chosen range of future times and market levels"[7].

Similar approaches can be implemented for down-and-in and down-and-out European calls. Actually, when the volatility is very large, these options can simply be replicated by some given quantities in the underlying and in a zero-coupon bond.

Carr & Chou (1996) extend the static hedging approach to a wide number of barrier options, such as: partial barrier options, forward starting barrier options, double barrier options, etc. Carr, Ellis and Gupta (1998) develop static hedges that allow for the "smile" or "frown" and which are based on the Put-Call symmetry, i.e.:

$$C(K)\frac{1}{\sqrt{K}} = P(H)\sqrt{H} \tag{3.10}$$

where

$$\sqrt{K\,H} = F$$

$C(K)$: call price with strike K

$P(H)$: put price with strike H

F: forward price

This relation is based on several assumptions: zero drift for the price process (zero cost of carry), no jumps and a certain symmetry condition for the volatility. Finally, the results are that a down-and-out call can be hedged by purchasing a standard call and selling off an instrument that has the same value that this call when the forward price reaches the barrier. Given the symmetry relation above, this instrument is just some quantity of a put. On the other hand, a down-and-in call can be hedged by purchasing some quantity of an out-of-the-money put, then, if and when the barrier is breached, the put is sold and a standard call is purchased instead. Similar hedges can be constructed for binary barrier options based on a more complicated Put-Call symmetry relationship.

[7] Derman, E., Kani, I., Ergener, D. (1994) - "Static Options Replication", The Journal of Derivatives, 2(4), 78-95

If the cost of carry is not zero, we may still find rather tight bounds for the replicating portfolio.

In general, barrier options seem to be easier to hedge statically than dynamically, but each hedging strategy relies on significant assumptions. So, finally, it is necessary to implement the strategy that best corresponds to a specific market.

Lookback options

Lookback options can also be hedged dynamically and this is apparently more efficient than for European calls given that the delta is smaller for the former than for the latter. However, given that gamma is, by contrast, higher, a dynamic trading strategy may still give rise to high transaction costs due to frequent rebalancing. The implementation is similar to the one for a call option and, most likely, the results are similar except for the case of transaction costs. Here, the performance of the hedge may be very different, i.e. smaller, due to the necessity of frequent rebalancing.

A most interesting fact is that if we use a stochastic volatility process to price and hedge these products, the results may be dangerously different form the Black-Scholes case. In particular, according to Davydov & Lintesky (2000), a CEV specification for volatility leads to very different deltas than under a log-normal model; actually, these new deltas my have different signs, so delta hedging under log-normality may produce extremely negative results. We may infer that for our Hull & White stochastic volatility model, the results may be similar.

A static or semi-static approach has also been looked for by different authors. Goldman, Sosin & Gatto (1979) price these options based on a hedging strategy. They propose for the case when the risk-neutral drift of the stock's process is 0 the followings: the floating strike lookback call can be replicated with a straddle that has an exercise price equal to the current minimum. This straddle is formed of an out-of-the-money put and an in-the-money call. Analogously, a floating strike lookback put is the same as a straddle formed by an in-the-money put and out-of-the-money call and an exercise price equal to the current maximum. At any point when the minimum (maximum) changes due to stock price variations, the straddle must be sold and another one formed.

Carr, Ellis & Gupta (1998) have produced a portfolio of standard options (calls and puts) that replicates such an option, based on the same Put-Call symmetry relationship. It may be that this portfolio undervalues the lookback option, but this disadvantage may be eliminated by increasing the number of options. Tompkins (2002) uses simulations to implement a very simple and intuitive “rolling-down strategy”. In particular, for a European floating strike lookback call, the first step is to buy a call with a strike price equal to the current minimum and the same maturity a s the lookback option. Then, provided that a new minimum price is reached, this call will be sold and the proceeds will be used to purchase another call with a strike price equal to the new minimum, and so on.

There are several super-replication approaches to hedging fixed strike lookback options such as the one proposed by Hobson (1998). The price is maximized under a specific model and the result represents the lowest initial investment needed for super-replicating the lookback option. Finally, two bounds on the lookback price are derived and it is shown that if the real price is outside these bounds, than an arbitrage profit can be generated.

Overall, it seems that lookback options are very sensitive to volatility, any hedging strategy should probably start by specifying a more realistic volatility model and use it accordingly.

Quanto options

Like exchange options, quantos will depend on the evolution of two assets: the underlying asset and the foreign exchange rate. Consequently, even though they may look as unaffected by the exchange rate’s evolution, the valuation and the hedging will be strongly affected by the correlation between the above processes. Once again, the dynamic hedging approach is viable and could be implemented. However, it is extremely complicated since it involves trading in the international markets, with funds being transferred permanently from one currency to another.

In particular, the premium of the option, in domestic currency, is deposited in a bank account. Then foreign currency is borrowed in order to implement an usual delta-hedging strategy. At each time step, the value of the foreign call is observed and, according to the fixed exchange rate, the value of the domestic currency account is modified. For example, if the value of the call has increased, the difference is calculated in domestic currency and borrowed domestically. Conversely, if the call value has decrease, the difference is also calculated in domestic currency and transferred from the domestic account in foreign currency.

In simple terms, as the value of the foreign call evolves over time, these transactions only insure that the domestic account contains the exact value of this call at the pre-established exchange rate. So, at the end, if the option will be out-of-the money, the proceeds in the account will be used to cover the borrowings in the foreign market. On the contrary, if the option will be in-the-money, the holdings in the domestic account are simply paid to the domestic investor. Then, the replicating portfolio is used to cover the initial borrowings for putting up the dynamic strategy.

Except for the rebalancing of the domestic account, this strategy is not different than the one implemented for a simple call. We can, thus infer, that the results will be similar. A dynamic approach is, for the time being, the only known possibility to hedge these products. However, it is also acknowledged that the correlation risk cannot be eliminated and this fact may have a strong impact on the hedge performance.

Digital options

Basically, digital options are bets on the asset being higher or below a pre-established level at expiration. It is commonly known that digitals are easy to price, but quite difficult to hedge because their delta increases to infinity close to the strike. Actually, in this region, the price of the digital option will resemble its delta, the delta will start to behave like gamma and the gamma will be closely related to the third derivative of the price with respect to the underlying asset. These facts explain why delta-hedging tends to be difficult to implement and may even break down in discrete intervals.

Under these circumstance, a static hedging approach has been adopted instead, consisting of replicating the digital with a call or put spread. For example, a short cash-or-nothing digital position can be replicated with a long bull call spread, that is by buying a call with a smaller strike and selling a call with a slightly higher strike. The closer the two strikes, the better the replication of the digital. The second call's strike is usually equal to the strike of the binary option. The difference between the two calls' strikes should be equal to the binary's option payoff. This way, the payoff is perfectly replicated with one exception: in the region between the two calls' strikes, the vertical spread will be worth more than the digital. Thus, such a spread is actually over-replicating the digital.

Ideally, the digital could be replicated exactly provided that the difference between the strikes of the two calls (or puts) tends to 0 and the quantity of spread needed tends to infinity. Realistically, such a strategy is subject to liquidity constraints in the options markets.

By simulations, Tompkins (2002) finds that this static approach is definitely better than the dynamic one since it saves transaction costs and performs well when the volatility is stochastic. Furthermore, Taleb (1997) suggests that: "the trader needs to shrink the difference between the strikes as time progresses until expiration, at a gradual pace. As such an optimal approach consumes transaction costs, there is a need for infrequent hedging."

Conclusions

The exotic options' utility is above all question and, as already shown, they can perform different functions for different end users. On the other hand, we must constantly be aware of the difficulties that these instruments pose. First, they are extremely hard to price and usually are very model dependent. Then, the opaque nature of OTC markets maintains the "exotisme" of these options, so that their risks are not always well-known. Careful hedging becomes a major issue and a difficult endeavor since the risks are more or less obscure.

In this paper, we have chosen a simulation approach to price and hedge exotic options. Due to its flexibility, Monte Carlo can accommodate a great variety of payoffs and the bounds it gives on prices are usually precise. We prove this assessment by constructing several Gauss programs for pricing path-dependent and correlation options. We compare the estimates obtained by 10000 simulations with the equivalent analytical formulas, most of them from Zhang (1998). For Asian options, lookback options, foreign-equity options and spread options the simulations are fairly accurate. However, for barrier options, the convergence is rather slow and the precision for quantos is ambiguous. For purpose of illustration, we have chosen these rather "simple" payoffs, but the main idea is that simulations are very efficient when the structures become very complicated and impossible to price analytically.

The previous results have been obtained in a pure Black-Scholes framework, so they rely heavily on a number of strong assumptions: the ability to trade continuously, no transaction costs, constant volatility, etc. We concentrate on the latter and review the attempts that have been made to model volatility deterministically or stochastically. The Hull & White (1987) model is chosen as a representative and implemented for the pricing of a standard European call, an arithmetic Asian option and a spread option. In parallel, the prices of these options are computed in a Black-Scholes framework, by simulations.

The chosen stochastic process for volatility leads to the same results for the European and Asian option: in a Black-Scholes environment, at-the-money options are overpriced while out-of and in-the-money ones are underpriced. The higher the volatility of volatility, the more pronounced this effect. However, the Asian option exhibits less overpricing (in absolute value) for at-the-money options and more severe underpricing for out-of and in-the-money ones than their standard counterparts. For spread options we observe that they seem to be constantly underpriced in a Black-Scholes world. However, the results are extremely sensitive to the parameters since two stochastic volatility processes are simulated, one for each underlying asset.

Next we address the sensitive subject of hedging exotic options and we relax the previously mentioned assumptions of the Black-Scholes world. First, we implement a dynamic hedging strategy for a European call and analyse the effect of discrete rebalancing, transaction costs (spreads and commissions on the traded amount) and stochastic volatility. Then, we construct a similar, but more complicated hedging procedure for an exchange option. We perform several comparative statistics by changing the frequency of rebalancing, the size of the commissions and the parameters of the stochastic volatility process in order to assess the effect of these imperfections on the hedging performance.

In a "perfect" market or in a market with no transaction costs, but with stochastic volatility, the mean hedging cost is approximately 0 for both options. However, the variability of the hedge is higher when stochastic volatility is introduced and also higher for the exotic option than for the standard one. The introduction of transaction costs significantly influences the hedge performance, particularly if combined with stochastic volatility. The effects are more dramatic for the exchange option than for a simple call since the rebalancing implies trading in two assets, not just one. Overall, it seems that the exotic option is far more sensitive to hedging errors than its simple counterpart.

Finally, we analyse and compare the dynamic and the static hedging approaches for several types of exotic options. The relative stability of the delta makes dynamic hedging very attractive for an Asian option. However, it is essential that this parameter be well-determined since it is easily affected by the observation frequency. The quality of the hedge is totally different for the multitude of barrier options: dynamic hedging is viable for "out" options if they are out-of-the-money close to the barrier and dangerous in other cases; static hedging appears as a better alternative for many of these options. Lookback options may pose problems due to their high gamma and they are subject to large errors when a different specification is used for the volatility process. For correlation options, in general, dynamic hedging is the only possibility by now, yet there may be problems related to the instability of the correlation coefficient.

In general, we show that Monte Carlo simulations are extremely useful for a variety of purposes, from pricing to hedging, form stochastic volatility modelling to risk analysis. In the field of exotic options, they have become a necessity.

Bibliography

[1] Ball, C. A., Roma, A. (1994) - "Stochastic Volatility Option Pricing", *The Journal of Financial and Quantitative Analysis*, 29, 589-607

[2] Bensaid, B., Lesne, J. P, Pages, H., Scheinkman, J. (1992) - "Derivative Asset Pricing with Transaction Costs", *Mathematical Finance* 2, 63-86

[3] Black, F., Scholes, M. (1973) - "The Pricing of Options and Other Corporate Liabilities", *Journal of Political Economy*

[4] Blattberg, R. C., and N. J. Gonedes (1974) - "A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices", *Journal of Business*, 47, 244-280

[5] Boyle, P., (1977) - "Options: A Monte Carlo Approach", *Journal of Financial Economics*, 4, 323-338

[6] Boyle, P., Broadie, M., Glasserman, P. (1997) - "Monte Carlo Methods for Security Pricing", *Journal of Economic Dynamics and Control*, 21, 1267-1321

[7] Boyle, P., Emanuel, D. (1980) - "Discretely Adjusted Option Hedges", *Journal of Financial Economics*, 8, 255-282

[8] Boyle, P., Vorst, T. (1992) - "Option Pricing in Discrete Time with Transaction Costs", *Journal of Finance*, 47, 271-293

[9] Bollerslev, T. (1986) - "Generalized Autoregressive Conditional Heteroscedasticity", *Journal of Econometrics*, 31

[10] Carr, P., Chou, A. (1996) - "Static Hedging of Barrier Securities", Working Paper

[11] Carr, P., Ellis, K., Gupta, V. (1998) - "Static Hedging of Exotic Options", *Journal of Finance*, 53, 1165-1190

[12] Castanias, R.P.(1979) - "Macroinformation and the Variability of Stock Market Prices", *Journal of Finance*, 34, 439-450

[13] Charnes, J. M. (2000) - "Using Simulation for Option Pricing", Working Paper, School of Business, The University of Kansas

[14] Clark, P. K. (1973) - "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices", *Econometrica*, 41,135-155.

[15] Das, S. R., Sundaram, R. K. (2000) - "A Discrete-Time Approach to Arbitrage-Free Pricing of Credit Derivatives", *Management Science*, 46, 46-62.

[16] Davis, M. (2003) - "Multi-Asset Options", Working Paper, Department of Mathematics, Imperial College, London

[17] Davydov, D., Lintesky, V., (2000) - "The Valuation and Hedging of Barrier Lookback Options for Alternative Stochastic Processes", Working Paper

[18] Dempster, A. H., Pliska, R.S. (2000) - "Mathematics of Derivative Securities", Cambridge University Press, New York

[19] Derman, E., Kani, I., Ergener, D. (1994) - "Static Options Replication", *The Journal of Derivatives*, 2(4), 78-95

[20] Dravid, A., Richardson, M., Sun, T. (1993) - "Pricing Foreign Index Contingent Claims: An Application to Nikkei Index Warrants", *Journal of Derivatives*, 1, 33-51

[21] Duan, J., Wei, J. (1999) - "Pricing Foreign-Currency and Cross-Currency Options under GARCH", *Journal of Derivatives*, 7(1)

[22] Edirisinghe, C., Naik, V., Uppal, R. (1993) - "Optimal Replication of Options with Transaction Costs and Trading Restrictions", *The Journal of Financial and Quantitative Analysis*, 117-138

[23] Engle, R. F. (1982) - "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation", *Econometrica*, 50

[24] Figlewski, S. (1989) - "Options Arbitrage in Imperfect Markets", *Journal of Finance* 44, 1289-1311

[25] Fouque, J. P., Papanicolaou, G., Sircar, K. R. (2000) - "Derivatives in Financial Markets with Stochastic Volatility", Cambridge University Press, New York

[26] Follmer, H, Schweizer, M. (1991) - "Hedging of Contingent under Incomplete Information", *Applied Stochastic Analysis*, 5, 389-414

[27] Follmer , H., Sondermann, D. (1986) - "Hedging of Non-Redundant Contingent Claims", *Contributions to Mathematical Economics*, 205-223

[28] Francois, P. (2001) - "A Note on Pricing and Hedging Asian Options with Observations Frequency", HEC Montreal, Working Paper

[29] Frey, R., (1997) - "Derivative Asset Analysis in Models with Level-Dependent and Stochastic Volatility", *CWI Quarterly*, 10, 1-34

[30] Galai, D. (1983) - "The Components of the Return from Hedging Options against Stocks", *The Journal of Business*, 56(1), 45-54

[31] Goldman, M. B., Sosin, H. B., Gatto, M. A. (1979) - "Path Dependent Options: Buy at the Low, Sell at the High", *The Journal of Finance*, 34(5), 1111-1127

[32] Goldman, B., Sosin, H., Shepp, L. A. (1979) - "On Contingent Claims that Insure Ex-Post Optimal Stock Marketing Timing", *Journal of Finance*, 34, 401-414

[33] Heinzl, T. A. (1999) - "Dynamic Hedging Strategies and Option Pricing in Bond Market Models with Transaction Costs", University of St. Gallen, Doctoral Dissertation

[34] Heston, S. (1993) - "A Closed Form Solution of Options Stochastic Volatility with Applications to Bond and Currency Options", *The Review of Financial Studies* 6, 327-343

[35] Ho, T. S., Stapleton, R. C, Subrahmanyam, M. G. (1995) - "Correlation Risk, Cross-Market Derivative Products and Portfolio Performance", *Journal of European Financial Managemen*t 1(2), 105-124

[36] Hobson, D. G., (1998) - "Robust Hedging of the Lookback Option", *Finance and Stochastics*, 2, 329-347

[37] Hull, J. C. (2000) - "Options, Futures and Other Derivatives", 5th edition, Prentice Hall, New Jersey

[38] Hull J., White, A. (1987) - "The Pricing of Options on Assets with Stochastic Volatility", *Journal of Finance*, 42, 281-300

[39] Johnson, H., Shanno, D. (1987) - "Option Pricing when the Variance Is Changing", *The Journal of Financial and Quantitative Analysis*, 22, 4, 143-151

[40] Kwok, Y., Wong, H. (1999) - "Currency-Translated Foreign Equity Options with Path-Dependent Features and Their Multi-asset Extensions", *International Journal of Theoretical and Applied Finance*, 3(2), 257-278

[41] Leland, H (1985) - "Option Pricing and Replication with Transaction Costs", *Journal of Finance*, 40, 1283-1303

[42] Levy, E. (1992) - "Pricing European Average Rate Currency Options", *Journal of International Money and Finance*, 11, 474-491

[43] Longstaff, F.A., Schwartz, E. S. (1995) - "Valuing Credit Derivatives", *The Journal of Fixed Income*, 6-12

[44] Kemna, A. G, Vorst, A. C. (1990) - "A Pricing Method for Options Based on Average Asset Values", *Journal of Banking and Finance*, 14, 113-129

[45] Margrabe, W. (1978) - "The Value of an Option to Exchange One Asset for Another", *Journal of Finance*, 33, 177-186

[46] Mercurio, F., Vorst, T. (1996) - "Option Pricing with Hedging at Fixed Trading Dates", *Applied Mathematical Finance*, 3, 135-158

[47] Merton, R. C. (1973) - "Theory of Rational Option Pricing", *The Bell Journal of Economics and Management Science* 4, 141-183

[48] Moessner, R. (2001) - "Over the Counter Interest Rate Options", Research Paper 1/2001, Bank of England, Centre for Central Banking Studies, London, England

[49] Mohamed, B. (1994) - "Simulations of Transaction Costs and Optimal Rehedging", *Applied Mathematical Finance*, 1, 49-62

[50] Rebonato, R. (1999) - "Volatility and Correlation: In the Pricing of Equity, FX and Interest-Rate Options", John Wiley & Sons Ltd.

[51] Reiner, E. (1992) - "Quanto Mechanics", *Risk* 5(3), 59-63

[52] Reiner, E., Rubinstein, M., (1991) - "Breaking Down the Barriers", *Risk* 4(8), 28-35

[53] Robins, R. P, Schachter, B. (1994) - "An Analysis of the Risk in Discretely Rebalanced Option Hedges and Delta-based Techniques", Working Paper

[54] Scott, L. O. (1987) - "Option Pricing when Variance Changes Randomly: Theory, Estimation, and an Application", *The Journal of Financial and Quantitative Analysis*, 22, 419-438

[55] Schweizer, M. (1995) - "On the Minimal Hedging Measure and the Fllmer-Schweizer Decomposition", *Applied Stochastic Analysis*, 13, 573-599

[56] Staum, J. (2002) - "Simulation in Financial Engineering", Working Paper, School of Operations Reasearch and Industrial Engineering, Cornell University

[57] Stein, E. M., Stein, C. J. (1991) - "Stock Price Distributions with Stochastic Volatility: An Analytic Approach", *Review of Financial Studies*, 4, 727-752

[58] Taleb, N. (1997) - "Dynamic Hedging, Managing Vanilla and Exotic Options", John Wiley & Sons Ltd.

[59] Toft, K. B., Reiner, E. S. (1997) - "Currency-Translated Foreign Equity Options: The American Case", *Advances in Futures and Options Research*, 9, 233-264

[60] Tompkins, R. G. (2002) - "Static versus Dynamic Hedging of Exotic Options: An Evaluation of Hedge Performance via Simulation", *The Journal of Risk and Finance*, 2002

[61] Turnbull, S., Wakeman, L. (1991) - "A Quick Algorithm for Pricing European Average Options", *Journal of Quantitative Analysis*, 26(3), 377-389

[62] Wilmott, P. (2001) - "Quantitative Finance", John Wiley & Sons, Ltd.

[63] Zhang, P. (1998) - "Exotic Options", 2nd edition, World Scientific Publishing.

We propose a new â€˜hedgedâ€™ Monte-Carlo (hmc) method to price -nancial derivatives, which allows to determine simultaneously the optimal hedge. The inclusion of the optimal hedging strategy allows one to reduce the nancial risk associated with option trading, and for the very same reason reduces considerably the variance of our hmc scheme as compared to previous methods.Â of the Monte-Carlo simulation. When minimizing the former by choosing the optimal strategy, we automatically reduce the latter.Â 4 American and other exotic options. The hmc method can be used to reduce the Monte Carlo error for all types of ex-otic options. We illustrate this point by showing how the method can be extended to price an American put option. â€¢ The advantage of the Monte Carlo simulation method is to deal with path dependent options. The superiority of the Monte Carlo simulation method is that it can simulate the underlying asset price path by path, calculate the payo associated with the information for each simulated path, e.g., Smax or Save, and utlize the average discounted payo to approximate the expected discounted payo, which is the value of path-dependent options. However, the advantage of the Monte Carlo simulation method causes the diculty to apply this method to pricing American options.Â Someone may try to apply the multiple-tier Monte Carlo simulation to estimating the holding value and thus price American options, but this method is infeasible for a large number of early exercise time points, n. Pricing and Hedging Asian Options. Where, S(ti) = Spot price at time t, N = number of equally distributed sample points T = time to maturity. In reality, most average price Asian options use arithmetic averaging over geometric averaging.Â 4.4 Monte Carlo Simulations: As defined in Exotic Option Trading, the principle of a Monte Carlo process is to generate a large number of finite paths, compute the payoff at each iteration, aggregate those payoffs, and subsequently divide that aggregated sum by the total number of simulated paths.