Université catholique de Louvain

Louvain School of Management (LSM)

Volatility Modelling in Option Pricing and its Impact on Payoff Replication Performance

Mémoire réalisé par
Nathan Lassance

Promoteur
Frédéric Vrins

Lecteur
Isabelle Platten

Année académique 2015-2016
Master en Ingénieur de gestion

"And finally, no matter how good the science gets, there are problems that inevitably depend on judgement, on art, on a feel for financial markets."

Martin Feldstein

Louvain School of Management

Summary

Master en Ingénieur de Gestion (INGE2)

Volatility Modelling in Option Pricing and its Impact on Payoff Replication Performance

Author: Nathan LASSANCE – NOMA: 3729 11 00

Supervisor: Frédéric VRINS

Volatility modelling in option pricing has been shown to be of first-order importance in improving upon the Black-Scholes pricing biases. However, no consensus emerges on its impact on hedging performance, with more realistic volatility specifications sometimes decreasing the accuracy compared to BS.

In light of this, we provide a thorough study of the hedging performance and behaviour of four models: constant volatility (Black-Scholes), local volatility (Practitioner Black-Scholes), GARCH-type (Heston-Nandi) and stochastic volatility (Heston). We apply a Delta-neutral payoff replication strategy on European call option quotes on the S&P500 and Apple.

We wish to assess the relative hedging performance of our models and how moneyness, maturity and volatility impact the accuracy. Additionally, we further dive into the link between hedging and pricing accuracy, the effect of a frequent parameters re-calibration and the drivers of under/over-hedging.

We use a pricing $MSE loss function on windows of two days and we calculate the hedging errors out-of-sample. Interesting results emerge:

- The maturity, moneyness and volatility have a substantial impact on hedging accuracy. A larger maturity sometimes surprisingly leads to lower errors for far-from-the-money options. Then, near-the-money options are better hedged than far-from-the-money ones under a low volatility, while the reverse effect is interestingly observed for long-term options under a high volatility. Finally, we observe that HN

and Heston are better suited to a more volatile asset: their errors are lower for Apple than the S&P500.

- More advanced volatility modelling overall decreases hedging accuracy. Under a low volatility, the ranking is always the same: PBS, BS, Heston, HN. When volatility rises, Heston and HN sometimes outperform BS and PBS for sufficiently OTM options. Based on the end-moneyness, the ranking shows that HN and Heston beat BS and PBS when the Delta is close to 0 during a large proportion of the hedging period.

- The worst models for hedging are overall the best ones for pricing, i.e. an over-fit of option prices leads to poor Delta estimates. Then, a frequent parameters re-calibration surprisingly increases the size of errors under a low volatility. When volatility rises, the ranking is modified for DOTM options, with HN and Heston beating BS and PBS. Finally, under/over-hedging is mainly driven by the re-balancing frequency: a lower frequency leads to more over-hedging.

Three implications emerge: (1) there is not one ideal model: the choice depends on the applied exercise, the underlying specificities, the option model parameters and the procedure followed; (2) adding complexity to make the procedure and models more realistic may counter-intuitively lead to worse results; (3) given the negative link between hedging and pricing performance, a focus should be put on application-oriented calibration procedures.

Acknowledgements

First and foremost, I am very grateful to my supervisor, Pr. Frédéric Vrins, for his availability and the very previous advice I got from each of our meeting. His expertise and research experience have been a source of inspiration for me throughout the writing of this master's thesis and confirmed my desire to dive further into academic research. I hope that this work will mark a first step towards fruitful future collaborations.

Besides my supervisor, my greatest thanks goes to my family. My mother, for her constant encouragement and the opportunities she enabled me to live during my studies and without whom I would not be the one I am today. My father, for the countless discussions we had together and who helped me nurture my curiosity towards the outside world. Finally, my brother, for the fun we had when we lived together and who never fails to compel me to think big in life.

I also dedicate these acknowledgements to my close friends who all hold a special place in my life: Matthieu, Vincent and Thomas who made these five years of studies a real blast; Maxime, Tanguy and François for the countless hours - sometimes fun, sometimes tiresome - spent with you on group projects; and finally a special thanks to my roommate Maxime for this last year lived with you and the delicious meals you never failed to share with me.

Last but not least, I thank you Chloé for what you have been bringing to my life since more than three years and for having tried to find some interest in my master's thesis in spite of the subject being far from what you like. Undoubtedly, all the breaks from our thesis spent together positively impacted this final version. You will never stop amazing me through your courage and your ability to always bring out the best in me.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Context of the Research

The appearance of financial markets shaped the modern economy and contributed to the economic growth by providing investment opportunities to a wide class of investors. Among them are derivative contracts, and options in particular. A derivative contract is "a financial instrument whose value depends on the values of other, more basic, underlying variables" (Hull, 2012). Options, in their plain vanilla versions, are a "financial instrument that gives one party the right, but not the obligation, to buy or sell an underlying asset from or to another party at a fixed price over a specific period of time" (Chance, 2003). In 2014, the total volume of futures and options contracts traded on exchanges worldwide was 21.87 billion (FIA, 2015). Globally, derivatives are estimated to represent from 10 to 20 times the size of the worldwide economy. In light of such numbers, we can not afford not to thoroughly understand how to properly model and manage them.

Options are used in the economy for various reasons: to profit from them, to speculate, to increase a portfolio performance and, more generally, to transfer risk from one party to another. We decided to focus on options in this thesis both due to their practical relevance and because they raise complex and interesting modelling challenges that have attracted a lot of research attention. Indeed, as this work will extensively show, options are highly non-linear and risky instruments. Consequently, pricing them and using them for risk management purposes must be dealt with great care and with sufficiently realistic models. It is therefore not surprising that the first option pricing model developed by Black and Scholes (1973) resulted in a Nobel Prize. However, the over-reliance on this same formula and more generally the deviances of complex derivatives products have been highly criticized during the recent financial crisis. Nowadays, regulatory (e.g. Basel, MiFID) and environmental pressures are pushing banks to upgrade their risk management.

In this context, we believe that developing and properly using realistic option pricing models is key for practitioners.

1.2 Research Questions & Motivation

In the literature, many different option pricing models have been developed, each with its pros and cons. We have chosen to analyze four different models. We carefully selected them based on their widespread reputation, theoretical features, empirical performance and practical relevance. This last point especially concerns the trade-off between complexity and computational intensity: a closed-form formula for the option price exists for each of the model considered. The four models are:

1. *The Black-Scholes (BS) model.* This is a continuous-time model that describes the underlying asset by a geometric Brownian motion with constant volatility.
2. *The Practitioner Black-Scholes (PBS) model.* This is a local volatility model that uses the BS formula with a deterministic volatility function as volatility input.
3. *The Heston-Nandi (HN) model.* This is a GARCH-type discrete-time model whose volatility is conditionally deterministic.
4. *The Heston model.* This is a continuous-time stochastic volatility model.

We limit ourselves to four models for matters of parsimony. This required to make some non-easy trade-offs. For example, we deliberately decided to set aside jump models even though they are crucial in pricing short-term options (Bakshi et al., 1997). Rather, we decided to focus on four distinct ways of modelling volatility.

Having outlined the context, the academic and practical relevance of the subject at hand and the models that we will evaluate, we define our main research question as:

> *Among the Black-Scholes, Practitioner Black-Scholes, Heston-Nandi and Heston models, which one performs best in terms of out-of-sample hedging performance, as measured by payoff replication, and how does performance differ according to the level of moneyness, maturity and volatility of the underlying?*

This research question raises several comments. We deliberately decided not to put our focus on pricing performance even though this measure is widely spread in the literature when comparing between different option pricing models. There are two main reasons for this decision:

- Pricing is in practice not so relevant for plain vanilla options that are typically well liquid and thus whose prices are driven by offer and demand. It is more relevant for exotic options or structured derivatives for which no liquid data are readily available.
- Although pricing is not practically relevant for plain vanilla options, it can still serve as a good measure to see how different models are able to capture information from option quotes, especially with out-of-sample forecasting. However, pricing performance is well documented in the literature (see Bakshi et al. (1997), Dumas et al. (1998), Heston and Nandi (2000) or Christoffersen and Jacobs (2004) among others), so that we do not think that it adds value to focus on this issue.

In practice, therefore, what matters for practitioners is rather whether these models are able to accurately hedge the position held in the option. We solely focus on this matter in our research. We assess the hedging performance out-of-sample and we use a non-traditional hedging criterion compared to the literature: payoff replication. In short, the Delta-neutral strategy can be used to try to replicate the payoff of the option at maturity. This reduces the uncertainty inherent to holding such a product, which is attractive for practitioners. Compared to the hedging tests in the literature, which are performed on an arguably small time period (five days in Baskhi et al. (1997), one week in Dumas et al. (1998) for instance), our hedging measure accounts for the whole life of the option and thus provides a more stringent evaluation criterion.

Our motivation and ambition with this research question arises from the fact that no consensus emerges from the literature on which model is best for hedging purposes, with even the simple BS model being able to outperform more realistic ones. Accordingly, we wish to contribute to a better understanding of the relative hedging performance of each model. Moreover, we wish to shed light on the drivers of hedging performance by classifying the errors according to moneyness and maturity and, importantly, by using two different time series having distinct levels of volatility: the S&P500 index (low volatility) and Apple (high volatility), whose options are very liquid. This contrasts with the tests in the literature that focus on a single underlying.

We will focus exclusively on European call options on equity. European-style options can only be exercised at expiration. This feature simplifies their analysis, and the models developed in this thesis present closed-form formulas for European options. Then, we focus only on call options because, if one assumes the put-call parity to hold true, then the dollar pricing errors are the same for calls and puts, hence making it unnecessary to delve into both types of contracts. Moreover, Bakshi et al. (1997), in their empirical study of the performance of alternative option pricing models on S&P500 data from 1988 to 1991, found the pricing and hedging results to be qualitatively similar for calls and puts. Finally, we opt for equity because this class of asset features well-known stylized facts that require more advanced option pricing models. In the literature, equity as underlying is also widely relied on. In practice, moreover, equity is often the riskiest portion of investors' portfolios (Bates, 2003) and options can be used to manage that risk.

In regards to our literature review and empirical findings, we identify three additional relevant sub-questions:

1. Some authors, e.g. Dumas et al. (1998) and Yung and Zhang (2003), point to the fact that, for hedging purposes, simpler models tend to fare better than more complex ones. However, these findings are not unanimous, see e.g. Petitjean and Moyaert (2011) who found that Heston outperforms BS in terms of hedging. With our two time series, we investigate this conundrum by comparing the out-of-sample payoff replication and pricing performance and ask ourselves the question:

   Is an accurate fit of the market option prices by a certain model a good indicator of its hedging performance?

   Indeed, one may argue that a model which accurately fits prices also accurately describes the underlying asset dynamics (if the model parameters are also reasonably stable through time). It should therefore also be good at hedging options on this same underlying. Put differently, we want to appraise whether the most accurate models for hedging are also as accurate for pricing, i.e. does an accurate fit of the option price curvature implies accurate Delta estimates?

2. Our procedure keeps the models' parameters constant until maturity to calculate the payoff replication errors. However, one could as well frequently re-estimate the parameters. We argue that re-calibrating the parameters sufficiently often may

change the ranking of the models performance by coping with the time instability of parameters for our more complex models. The question therefore is:

What is the effect of a frequent re-calibration of parameters on the amplitude of payoff replication errors and on the relative performance of our models compared to keeping parameters constant over the life of option?

3. Our empirical results point to a systematic over-hedging of the payoff at maturity. From Joshi (2008), we argue that this is due to the positive Gamma effect, and therefore that the re-balancing frequency, the underlying asset expected return and the underlying asset volatility may all have an impact on whether we observe an under or over-hedging. We assess this theoretically, i.e. on simulated data, because this gives us flexibility in the variables whose effect we want to test. Our payoff replication procedure proves useful to that matter. Indeed, in the contrary to traditional hedging tests that are performed only empirically, we can simulate the replicating portfolio until maturity under our different model price dynamics. We rely on the HN model and our specific question is:

   How does under/over-hedging depend on the re-balancing frequency, the underlying asset expected return and the underlying asset volatility?

1.3 Summary of Results

Based on the S&P500 (low volatility) and Apple (high volatility) time series, the main results uncovered through our study are stated below. Because the study is based on two time series, we do not pretend that the findings can automatically be generalized, especially since the two time series do not only distinguish by their volatility (Apple presents more negative skewness and leptokurticity). Our goal is rather to get deep insights into the hedging behaviour and performance of different classes of models and to compare our findings with the literature.

Effect of maturity. The payoff replication absolute errors increase with maturity for OTM and ITM options. However, the errors sometimes decrease with maturity for DOTM and DITM options. This is counter-intuitive since we would except payoff replication to be less accurate when you have to hedge during a longer period of time.

Effect of moneyness (at inception). Under a low volatility, we often observe U-shaped curves, meaning that that near-the-money options are more accurately hedged than far-from-the-money ones. When the volatility is higher, we strikingly observe the reverse effect for long-term options.

Models relative performance. More advanced volatility modellings generally lead to less accurate hedging. When the volatility is low, the ranking is always the same: 1. PBS, 2. BS, 3. Heston, 4. HN. When volatility rises, Heston and HN perform very closely and outperform BS and PBS for long-term OTM and DOTM options. For ITM and DITM options, BS and PBS perform best and display close results. A look at the ranking in function of the end-moneyness (at maturity) shows that HN and Heston achieve a more accurate hedge than BS and PBS when the Delta is close to 0 during a large enough proportion of the hedging period.

Effect of volatility. More stable underlying asset returns do not always induce more accurate hedging. Specifically, a lower underlying volatility (S&P500) overall leads to lower absolute payoff replication errors, measured in percentage of the option premium, for BS and PBS, and the contrary for HN and Heston. Lower errors for a higher volatility (Apple) do interestingly emerge under BS and PBS for DITM options.

Link with pricing performance. The worst models for hedging are overall the best models for pricing. Indeed, HN and Heston are overall performing best, with PBS still doing quite well. This indicates that an over-fit of option prices generates poor Delta estimates and that a hedging-based estimation function should ideally be applied rather than the pricing \$MSE standard practice.

Effect of a frequent re-calibration. Under a low volatility, Heston and HN are still outperformed and the errors are surprisingly nearly always larger than without re-calibration: time-varying parameters do not generate more accurate Delta estimates. When volatility rises, HN and Heston outperform BS and PBS for DOTM options, in the contrary to the case of constant parameters. Re-calibrating is now frequently generates lower errors, especially for DOTM options.

Drivers of under/over-hedging. We observe a systematic over-hedging on our empirical data, and very strikingly under a low volatility. Based on Monte Carlo simulation of the HN model, we show that the main driver is the re-balancing frequency: a low frequency leads to a high proportion of over-hedging, and vice versa. This is due to the option price convexity, which in turn implies that a higher maturity leads to more over-hedging. Additionally, the errors amplitude is larger under a low frequency because larger movements in the underlying asset can happen between two revision dates.

1.4 Structure of the Master's Thesis

Having introduced the context of our work, the research questions and our main results, the rest of the master's thesis is structured along the following chapters.

Chapter 2 We begin by reviewing the concepts that are necessary to grasp the remaining of the thesis. We cover elements on options, Delta hedging, the stylized facts of stock returns, stochastic calculus, risk-neutral valuation, the important concepts of no-arbitrage and market completeness, and end up with Monte Carlo simulation.

Chapter 3 This chapter covers the BS model. Based on the model assumptions, we derive the BS formula and present the payoff replication strategy, which is a central concept for our thesis. We then introduce the Greeks of the BS model. Finally, we review the model limitations and the concepts of volatility smile and skew.

Chapter 4 In this chapter, we review the three alternative models used in this thesis. The PBS model is first presented, along with its empirical performance. The GARCH and stochastic volatility frameworks are then introduced along the same structure: theoretical framework, model choice rationale, main features, implementation, Delta formula and finally empirical performance resurging from the literature.

Chapter 5 Here, we detail the methodology that we followed to answer to our research question and sub-questions. We describe the data that we used and how we filtered them. Then, we turn to the empirical study of payoff replication where we explain the market calibration procedure, our sliding window technique and the error measures used to assess the performance of the models. We end up by explaining how we will assess the over-hedging effect based on Monte Carlo HN simulations.

Chapter 6 In this chapter, we detail our results as summarized here before and discuss the outcomes. Prior to that, we provide statistical descriptions of the two time series and a short analysis of parameters. The analyses are performed on Matlab. The main Matlab functions that we used can be found on the CD attached to this master's thesis, while appendix A describes each of them.

Chapter 7 This final chapter concludes the thesis. We point out implications and underline the limitations of our work to finish with some suggestions for future research.

Chapter 2

Option Pricing: Concepts & Tools

2.1 Properties of Options

Options are part of one of the most prominent categories of financial assets, derivatives, that we defined in the introduction. Options themselves include different instruments. They can be divided between plain vanilla and exotic options. Exotic options are more complex in nature and will not be covered in this thesis. Among plain vanilla options, we distinguish between call and put options. A call option grants the right to buy the underlying, while a put option grants the right to sell the underlying. From now on, we will refer to the underlying as being a stock and we denote its price at time t as S_t. The fixed price at which you can buy or sell the stock is the strike price that we denote K. Options also have an expiration date, and therefore a certain maturity $T - t$ at time t. An option can be European, if it can only be exercised at expiration, or American, if it can be exercised at any time before expiration. We recall that we will solely focus on European options in this thesis. One other important characteristic of an option is its payoff, which is the option value at expiration:

$$\begin{aligned} C_T &= \max(S_T - K, 0) \\ P_T &= \max(K - S_T, 0) \end{aligned} \tag{2.1}$$

Finally, one important result in option theory is the put-call parity, which links the value of a call and a put option (with the same strike and maturity) as follows:

$$P_t + S_t = C_t + Ke^{-r(T-t)} \tag{2.2}$$

Given that this relation holds well for very liquid options such as the one analyzed in the empirical part, we restrict our investigation to call options from now on.

2.2 Delta Hedging

Options are leveraged instruments: they can be used to amplify gains, but also losses. However, most investors do not opt for speculative strategies. Rather, options are used for hedging purposes. Hedging has a specific meaning in finance which is "risk reduction with offsetting transactions that usually involve derivative securities" (van der Wijst, 2013). The most common hedging strategy used by traders is *Delta hedging* (Yung and Zhang, 2003). Delta hedging caps the total amount that can be lost when holding a certain number of stocks, by selling call options on this stock. This strategy can be used, theoretically, to create a risk-free portfolio. Indeed, suppose you are short a call option which is a certain function C_t that depends on time, the stock price and other parameters as we will see when introducing the BS model in chapter 3. Then, we can define the Delta of the option as the derivative with respect to the stock price:

$$\Delta_t \equiv \frac{\partial C_t}{\partial S_t} \tag{2.3}$$

By holding a portfolio composed of shorting one call option C_t and buying an amount of stocks given by Δ_t, we create a hedged portfolio: we talk about a *Delta-neutral portfolio*. However, the Delta does not remain constant over time (the second-order derivative, Gamma, must be accounted for). This means that the portfolio will theoretically remain hedged for only an infinitesimal amount of time. In practice, the amount of stock held will thus have to be adjusted periodically (e.g. daily), which is known as *re-balancing* or *dynamic hedging* (Hull, 2012). However, assuming that re-balancing can be done continuously, we can theoretically obtain a risk-free portfolio, which provides an argument for finding the fair (arbitrage-free) price of the option (Joshi, 2008). This is the BS price of the option, that we will derive in chapter 3. This means, as section 3.3 will show, that investing the BS price of the option in this continuously re-balanced Delta-neutral portfolio will perfectly replicate the option payoff under the model assumptions: this is called a *payoff replication strategy*. This is desirable for institutions, which would therefore be able to replicate the payoff at maturity that they will be liable to pay to investors who bought options from them. More generally, the economic attractiveness of hedging rests mainly "on the smoother, less volatile performance that hedging brings about" (van der Wijst, 2013). Hedging performance and behaviour, through payoff replication, will be the subject of our empirical study.

2.3 Stylized Facts of Stock Returns

Stock returns do depict so-called *stylized facts* that make their distribution differ from the Gaussian distribution. We describe three characteristics here: (1) kurtosis, skewness and leverage effect, (2) volatility clustering and (3) the mean-reverting behavior of volatility. These properties have to be adequately modeled in order to achieve an accurate stock price dynamic and thus develop a well-performing option pricing model.

2.3.1 Skewness, Kurtosis & Leverage Effect

The skewness (S) and kurtosis (κ) are the third and fourth standardized moments of a probability distribution X:

$$S = E\left[\left(\frac{X-\mu}{\sigma}\right)^3\right] \;\; ; \;\; \kappa = E\left[\left(\frac{X-\mu}{\sigma}\right)^4\right] \tag{2.4}$$

The skewness measures the degree of asymmetry of the distribution, while the kurtosis measures the tailedness of the distribution. The Gaussian distribution has values of $S = 0$ and $\kappa = 3$. In the case of stock returns, we typically have $S < 0$, i.e. a long tail to the left and so more downside risk than the Gaussian distribution, and $\kappa > 3$, i.e. fatter tails than the Gaussian. A negative skewness means that there is an asymmetric relation between returns and volatility: negative returns impact volatility relatively more than positive ones. The reason often mentioned is the so-called *leverage effect* of Black (1976): as the stock price drops, the leverage of the firm increases, the risk increases and hence the volatility raises too. Section 3.5 will explain that these deviations from normality give rise to the so-called *volatility smile* of the BS model.

2.3.2 Volatility Clustering

Volatility clustering refers to fact that, as noted by Mandelbrot (1963), the volatility of stock returns is serially correlated: large movements tend to be followed by large movements, and similarly for small movements. This results in volatility clusters where high and low volatility periods are grouped together. Moreover, the auto-correlation function tends to be very persistent, i.e. slowly decaying. This serial correlation of volatility is at the origin of the family of ARCH and GARCH models that have also been applied to option pricing, as we will analyze in section 4.2.

2.3.3 Mean-Reverting Volatility

When modelling volatility, we often opt for a mean-reverting process, meaning that the instantaneous volatility reverts to its long-term value at a certain rate. This property makes economical sense since volatility cannot increase or decrease indefinitely but has rather a natural level which is occasionally perturbed (Joshi, 2008). This effect is verified in practice. For example, Engle and Patton (2001) explain that option prices are consistent with mean reversion: implied volatilities (see section 3.5) of long-maturity options are less volatile and close to the long run average volatility of the underlying.

2.4 Stochastic Tools for Option Pricing

In this thesis, except for the HN model, we will deal with continuous-time models. These models are based on stochastic processes and stochastic calculus, which differ from their deterministic counterparts. In this section, based on Brigo and Mercurio (2006), Hull (2012) and van der Wijst (2013), we report the main tools needed to properly understand the models analyzed further.

2.4.1 Stochastic Differential Equations

The class of processes with which we will work are stochastic differential equations (SDE), which are of the form:

$$dS_t = \mu(t, S_t)dt + \sigma(t, S_t)dW_t \tag{2.5}$$

with $S_0 > 0$. The function $\mu(t, S_t)$ is the deterministic part of the process called the *drift*, and $\sigma(t, S_t)$ the *diffusion coefficient*. The randomness in this equation enters from the term dW_t and the initial condition S_0.

The process W_t is called a *Brownian motion* and is defined by three characteristics for any $0 < s < t < u$ and any $h > 0$: (1) Independent increments: $W_u - W_t \perp\!\!\!\perp W_t - W_s$, (2) Stationary increments: $W_{t+h} - W_{s+h} \sim W_t - W_s$ and (3) Gaussian increments

$W_t - W_s \sim \mathcal{N}(0, t-s)$. The paths $t \mapsto W_t$ are required to be continuous, and $W_0 = 0$. These properties imply that W_t is nowhere differentiable, which is why we must have recourse to stochastic calculus. Finally, one special case of equation (2.5) is when $\mu(t, S_t)$ and $\sigma(t, S_t)$ are directly proportional to S_t. This is called a *geometric Brownian motion* (GBM), which is the price dynamic used in the BS model. It is of the form:

$$dS_t = \mu S_t dt + \sigma S_t dW_t \tag{2.6}$$

Figure 2.1 displays an example of a simulation of the GBM process.

2.4.2 Itô's Lemma

One very important stochastic calculus result in finance is the well-known Itô's lemma, which gives the corresponding chain-rule for differentials in a stochastic context. This is based on three useful stochastic formulas: $dW_t dW_t = dt$, $dtdt = 0$ and $dW_t dt = 0$.

Suppose that we have a SDE as given in equation (2.5). Then, given a smooth function $\phi(t, S_t)$, Itô's lemma says that the function ϕ follows the process

$$d\phi(t, S_t) = \left(\frac{\partial \phi}{\partial S_t}\mu(t, S_t) + \frac{\partial \phi}{\partial t} + \frac{1}{2}\frac{\partial^2 \phi}{\partial S_t^2}\sigma^2(t, S_t)\right)dt + \frac{\partial \phi}{\partial S_t}\sigma(t, S_t)dW_t \tag{2.7}$$

where, importantly, the Brownian motions in equations (2.5) and (2.7) are the same.

2.4.3 Log-Normal Distribution & Martingale

If we apply Itô's lemma to $\ln S_t$ where S_t follows a GBM process as given in equation (2.6), we have the following process:

$$d\ln S_t = \left(\mu - \frac{\sigma^2}{2}\right)dt + \sigma dW_t \tag{2.8}$$

whose solution at some future time T is, with $W_T \sim \mathcal{N}(0, T)$,

$$S_T = S_0 e^{\left(\mu - \frac{\sigma^2}{2}\right)T + \sigma W_T} \tag{2.9}$$

FIGURE 2.1: Stock price simulation under the GBM process.

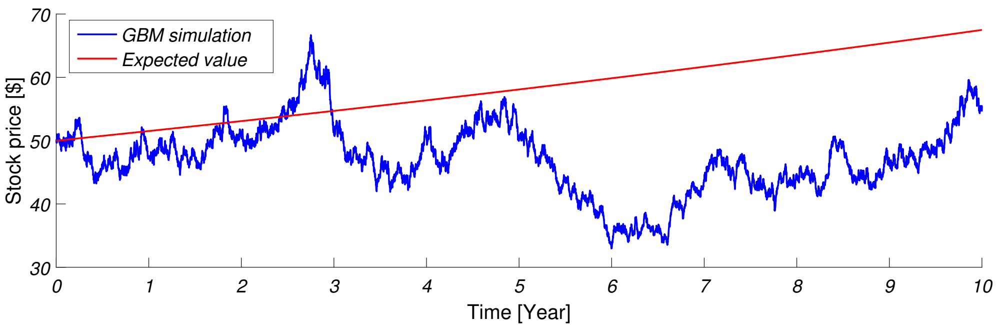

Notes. The simulation was done under the physical measure with a Milstein discretization scheme[1] and is compared with the process expected value. Parameters' values: $S_0 = 50$, $\mu = 5\%$, $\sigma = 20\%$, $T = 10$. Number of time steps: 10000.

This means that S_T follows a log-normal distribution since equation (2.9) implies

$$\ln S_T \sim \mathcal{N}\left[\ln S_0 + \left(\mu - \frac{\sigma^2}{2}\right)T,\ \sigma^2 T\right] \tag{2.10}$$

Figure 2.2 shows the density function of stock returns under a GBM process obtained by Monte Carlo simulation (see section 2.7). As expected, the distribution is very close to the normal density (since the log-price is log-normally distributed).

Finally, if we define a new process $Y_t = e^{-\mu t} S_t$, then by Itô's lemma we obtain $dY_t = \sigma Y_t dW_t$. Since this process has a zero drift term, we have that $e^{-\mu t} S_t$ is a *martingale*. A martingale is a process such that, for all $t < T$, we have:

$$\mathbb{E}[X_T | \mathcal{F}_t] = X_t \tag{2.11}$$

where $\mathcal{F}_t$ is the information set until date t (called filtration). This property "is a picture of a "fair game", where it is not possible to gain or lose on average" (Brigo and Mercurio, 2006). This is a crucial property to model the absence of arbitrage in finance (more on that in section 2.6). This is also extensively used for derivatives pricing: we price the

[1]To simulate continuous-time processes for Monte Carlo purposes (see section 2.7 for an explanation of Monte Carlo simulation), we need to discretize the processes (since Monte Carlo is done in discrete time steps). Many discretization schemes exist to achieve this, and in this thesis we apply the Milstein scheme (see Gatheral (2006) and Rouah (n.a.) for more details). The Milstein discretization is obtained by a second-order Itô-Taylor expansion of the process and, compared to the more simple single order Euler scheme, is not more computationally intensive and converges more rapidly. For the BS model, the Milstein discretization of the stock price process under the physical measure is: $S_{t+\Delta t} = S_t + \mu S_t \Delta t + \sigma S_t Z \sqrt{\Delta t} + \frac{1}{2}\sigma^2 \Delta t (Z^2 - 1)$ where $Z \sim \mathcal{N}(0, 1)$ and Δt is the discrete time step.

[2]As opposed to risk-neutral, see section 2.5.

FIGURE 2.2: Physical density of stock returns under the GBM process.

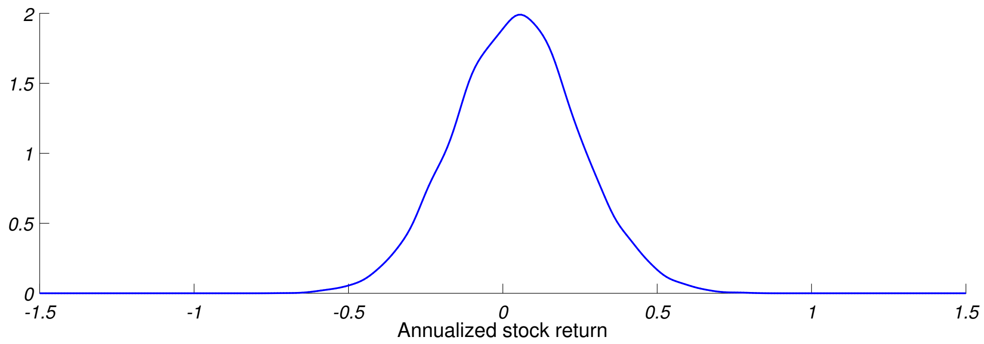

Notes. The physical[2]density was obtained by Monte Carlo simulation, with a Gaussian kernel under Matlab default bandwidth. The simulation was done with a Milstein discretization scheme for $R_t \equiv \frac{dS_t}{S_t} = \mu dt + \sigma dW_t$. Parameters' values: $\mu = 5\%$, $\sigma = 20\%$, $T = 1$. Number of time steps: 1000. Number of Monte Carlo random paths: 10000.

option under a probability measure chosen such that the price is simply given by the discounted expected value of the payoff under the risk-free rate. This probability measure is not the physical (i.e. real) measure, but the risk-neutral one.

2.5 Equivalent Martingale Measure & Risk-Neutral Valuation

This section is based on Brigo and Mercurio (2006). Option pricing relies on an equivalent martingale measure method. The principle is to change the expectation operator in the calculation of the option price from the physical measure $\mathbb{P}$ to another equivalent measure $\mathbb{Q}$, under which the properly discounted asset price is again a martingale. Two measures $\mathbb{P}$ and $\mathbb{Q}$ are said to be equivalent if "they share the same sets of null probability (or of probability one, which is equivalent)." What is interesting is that properly changing the probability measure changes the drift of the process, but keeps the same diffusion coefficient. Mathematically, when two measures are said to be equivalent, we can express the first in terms of the second in the following way:

$$\left.\frac{d\mathbb{Q}}{d\mathbb{P}}\right|_{\mathcal{F}_t} = \rho_t \tag{2.12}$$

where ρ_t is a martingale called the Radon-Nidokym derivative. Under Girsanov theorem, and considering the general process $dS_t = \mu(S_t)dt + \sigma(S_t)dW_t$, the change of measure is such that the new Brownian motion $W_t^{\mathbb{Q}}$ is given by

$$dW_t^{\mathbb{Q}} = -\left[\frac{\mu^{\mathbb{Q}}(S_t) - \mu(S_t)}{\sigma(S_t)}\right] dt + dW_t \tag{2.13}$$

where $\mu^{\mathbb{Q}}(S_t)$ is the drift under the new probability measure. In option pricing theory, we set $\mu^{\mathbb{Q}}(S_t) = r$, so that the probability measure $\mathbb{Q}$ is called the risk-neutral measure. In this way, the Brownian motion incorporates the market price of risk:

$$dW_t^{\mathbb{Q}} = \left[\frac{\mu(S_t) - r}{\sigma(S_t)}\right] dt + dW_t \tag{2.14}$$

As a result, the stochastic process for the stock price under the risk-neutral measure is

$$dS_t = rdt + \sigma(S_t) dW_t^{\mathbb{Q}} \tag{2.15}$$

under which we have

$$e^{-r(T-t)}\mathbb{E}^{\mathbb{Q}}[S_T|\mathcal{F}_t] = S_t \tag{2.16}$$

i.e. we have our new equivalent martingale measure. This new process can be used to derive the BS formula of an option without having to know the value of the risk premium of the underlying stock (see Cox and Ross, 1976). Next chapter will derive the formula by another method, but leading to the same conclusion: the BS price of an option does not depend on the underlying expected return μ. Risk-neutral valuation is also at the basis of the other more advanced models presented in this thesis: the physical processes that describe the underlying dynamics always have to be transformed in the risk-neutral processes to find the option price and Delta formulas.

2.6 No-arbitrage & Market Completeness

The previous section introduced the concept of equivalent martingale measure used to price options. As explained by Brigo and Mercurio (2006), the existence and uniqueness of the equivalent martingale measure is linked to the arbitrage-free and market completeness concepts. Specifically, we have that (1) the market is arbitrage-free if and only if there exists an equivalent martingale measure, (2) the market is complete if and only if the equivalent martingale measure is unique.

Absence of arbitrage refers to the impossibility to invest zero today and receive a non-negative amount in the future that is positive with a positive probability. It implies that two portfolios having the same payoffs at a given future date must have the same price today. This property will be respected in all the models used in this thesis.

Market completeness means that any contingent claim is attainable. In other words, every payoff can be perfectly replicated via a self-financing strategy[3] (Ortega, 2009). In particular, this holds true for the payoff at maturity of an option. Market completeness is verified by the BS model, but not by the other three models.

2.7 Monte Carlo Simulation

To conclude this chapter, we present a simulation technique extensively used in finance: *Monte Carlo simulation*. This can be used whenever one wants to compute a quantity at time $T > t$ dependent on the value of a certain path-dependent process.

To illustrate, suppose we want to find the value of a call option under a GBM process. The most straightforward way is to use the BS formula. Alternatively, although unnecessarily more computationally intensive, we can use Monte Carlo simulation. In that case, we start from the following risk-neutral process: $dS_t = rS_t dt + \sigma S_t dW_t$. Using discretization schemes (as explained on figure 2.1), one can discretize this continuous process to simulate it and find the value of a call option under Monte Carlo simulation. The procedure goes as follows: a set of MC random paths[4] $(\varepsilon_{t+1,j}, ..., \varepsilon_{T,j}), j = 1, ..., MC$ is generated (with $\varepsilon_t \sim \mathcal{N}(0, 1)$) and the corresponding prices $S_{T,j}$ are calculated. The price C_t of the call option is then approximated by

$$\hat{C}_t = e^{-r(T-t)} \frac{1}{MC} \sum_{j=1}^{MC} \max(S_{T,j} - K, 0) \tag{2.17}$$

We will use this very handy technique to generate probability density functions of stock returns from our models, as well as in the empirical part to analyze the over-hedging behaviour based on simulated data from the Heston-Nandi stock price dynamic.

[3] A self-financing strategy is "an investment strategy that does not call for external cash-flows" (Vrins, 2015).

[4] Around 5000 paths are sufficient to get a reasonably precise estimate (Christoffersen, 2012).

Chapter 3

Black-Scholes Framework

One of the greatest breakthroughs in option pricing and in the theory of finance has been the paper of Black and Scholes (1973), developed with the help of Robert Merton. The paper derived a new method to determine the fair value of derivatives. In this chapter, we derive what is now known as the *Black-Scholes formula*, we present the payoff replication strategy and finally we discuss the BS empirical limitations justifying more advanced models that will be described in chapter 4.

3.1 Model Assumptions

From Black and Scholes (1973) and Hull (2012), we can hold these 9 assumptions:

1. The underlying follows a GBM process as in equation (2.6), with σ being constant.
2. The risk-free interest rate r is constant.
3. The underlying asset involves no dividends or other cash flow distributions.
4. The option is of European-style.
5. There are no transaction costs or taxes involved.
6. There are no riskless arbitrage opportunities.
7. All securities are perfectly divisible.
8. Security trading is continuous and not discrete.
9. The short selling of securities with full use of proceeds is permitted.

In this thesis, we will relax three assumptions. First, the BS formula and the other option pricing models can easily be adapted to securities involving cash flow distributions. Second, the Heston-Nandi model of section 4.2 is a discrete-time model. Finally, and most importantly, the assumption of a constant volatility of asset returns goes strongly against empirical evidence. The resulting limitations will be discussed in section 3.5 and the alternative models in chapter 4 will try to model the volatility more realistically.

3.2 Black-Scholes PDE & Formula

3.2.1 Derivation of the Formula

There are several ways by which we can derive the famous BS formula. One possibility is by risk-neutral valuation as hinted in section 2.5, which is formally proven in van der Wijst (2013). In this section, we rely on the original proof of Black and Scholes. Their derivation gives insights into how constructing a Delta-neutral portfolio can remove uncertainty and therefore lead to the fair price of an option. The elements of proof of this section rely on Black and Scholes (1973), Hull (2012) and van der Wijst (2013).

We assume that the stock price follows a GBM process, given by

$$dS_t = \mu S_t dt + \sigma S_t dW_t \tag{3.1}$$

If we refer to f_t as the price of the derivative on S_t, then, from Itô's lemma, we have:

$$df_t = \left(\frac{\partial f_t}{\partial S_t} \mu S_t + \frac{\partial f_t}{\partial t} + \frac{1}{2} \frac{\partial^2 f_t}{\partial S_t^2} \sigma^2 S_t^2 \right) dt + \frac{\partial f_t}{\partial S_t} \sigma S_t dW_t \tag{3.2}$$

Since the Brownian motion in equation (3.2) is the same as the one from equation (3.1), we can make a portfolio of the stock and the derivative from which the Brownian motions disappear, i.e. a risk-free portfolio. This portfolio, that we denote as Π_t, is the following: $\Pi_t = -f_t + \Delta_t S_t$. In this equation, Δ_t is the option Delta, i.e. $\frac{\partial f_t}{\partial S_t}$. The dynamics of the portfolio is therefore given by[1]

$$d\Pi_t = -df_t + \Delta_t dS_t \tag{3.3}$$

By substituting (3.1) and (3.2) into (3.3), we get

$$d\Pi_t = -\left(\frac{\partial f_t}{\partial t} + \frac{1}{2} \frac{\partial^2 f_t}{\partial S_t^2} \sigma^2 S_t^2 \right) dt \tag{3.4}$$

[1] This formula is a shortcut, since we have $d(\Delta_t S_t) = \Delta_t dS_t + d\Delta_t S_t + d\Delta_t dS_t$. However, we can prove that the derivation is still true with the *self-financing condition* (Rosu and Stroock, 2004).

We see that the process Π_t contains no diffusion term. By the assumption of no arbitrage, Π_t thus has a return exactly equal to the risk-free rate: $d\Pi_t = r\Pi_t dt$. Substituting equation (3.4) and the definition of the portfolio Π_t into this last equation, we get:

$$\frac{\partial f_t}{\partial t} + \frac{1}{2}\frac{\partial^2 f_t}{\partial S_t^2}\sigma^2 S_t^2 + \frac{\partial f_t}{\partial S_t} S_t r - r f_t = 0 \tag{3.5}$$

This is the BS partial differential equation (PDE). It holds for any derivative satisfying the assumptions of section 3.1. We see that this equation does not depend on the expected return of the stock μ. Joshi (2008) clearly explains why this is the case in the BS model: "[...] in the Black-Scholes model, we can eliminate all the risk by continuous trading. Once the risk has been eliminated, there can be no risk premium."

To solve this PDE, the boundary conditions must be specified. For a call option, they are: $f_t = 0$ if $S_t = 0, \forall t$ and $f_T = \max(S_T - K, 0)$. The equation (3.5) with these boundary conditions has a unique solution, which is the BS formula:

$$C_t = S_t \mathcal{N}(d_1) - K e^{-r(T-t)} \mathcal{N}(d_2) \tag{3.6}$$

where

$$\begin{aligned} d_1 &= \frac{\ln(S_t/K) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}} \\ d_2 &= \frac{\ln(S_t/K) + (r - \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}} \end{aligned} \tag{3.7}$$

From Gatheral (2006), $\mathcal{N}(d_1)$ and $\mathcal{N}(d_2)$ are risk-adjusted pseudo-probabilities: $S_t\mathcal{N}(d_1)$ is the pseudo-expectation of the final stock price given that the option is in-the-money and $\mathcal{N}(d_2)$ is the pseudo-probability that the option will be exercised.

Following Lehar et al. (2002), we use option quotes to estimate the BS model implied volatility (see section 3.5) σ^* yielding the lowest pricing error. Following Christoffersen and Jacobs (2003), who showed that the \$MSE estimates perform best across different loss functions, we estimate σ^* by non-linear least squares as follows:[2]

$$\sigma^* = Arg \min_{\sigma} \$\text{MSE} \equiv Arg \min_{\sigma} \frac{1}{n} \sum_{i=1}^{n} \left(C_i^{mkt} - C_i^{BS}(S_t, r_i, K_i, T_i; \sigma) \right)^2 \tag{3.8}$$

where n is the number of observed option prices in the sample considered.

[2]This type of procedure is called calibration to market data.

3.2.2 Adaptation to Cash Flow Distribution

The BS formula can easily be adapted to the case where the underlying asset involves some sort of cash flow distribution. Following the literature practice, we adjust the stock price by the present value of dividends between t and the expiration date T (see Baskhi et al. (1997), Dumas et al. (1998) or Yung and Zhang (2003) for example).

3.3 Market Completeness & Payoff Replication Strategy

As mentioned in section 2.6, one of the most important properties of the BS model is market completeness. This property implies that the BS price of an option is such that if it is invested in a continuously re-balanced Delta-neutral portfolio, the option payoff at maturity will always be perfectly replicated under the model assumptions. This is called the *replication strategy* of the option's payoff. The replicating portfolio Π_t follows the following process (Vrins, 2015):

$$d\Pi_t = \Delta_t dS_t + r(\Pi_t - \Delta_t S_t)dt \tag{3.9}$$

with $\Pi_0 = C_0$ (the option premium). This equation means that we adjust the amount of stock held according to the value of the Delta, borrowing money at a rate r if the position in the stock needs to be increased, and placing money at that rate otherwise. The replicating portfolio is constructed with only two tradable assets: the underlying stock and a cash account earning the risk-free rate r.

Figure 3.1 illustrates the replication strategy in the BS framework. We see that the replicating portfolio converges to the payoff, and is *exactly equal* to the payoff at maturity. However, the assumptions of the BS model are not verified in reality and continuous re-balancing is not practically feasible. Therefore, one of the goals of this thesis will be to assess how accurate is BS for payoff replication when using empirical data. In brief, despite its shortcomings, we will show that BS compares overall very favourably to Heston-Nandi and Heston but is outperformed by the PBS model.

FIGURE 3.1: Payoff replication strategy under the GBM process.

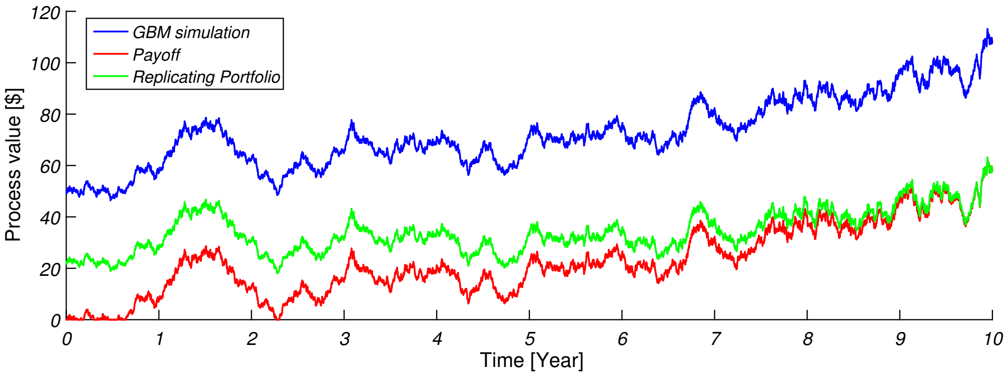

Notes. The simulation was done under the physical measure with a Milstein discretization scheme. Parameters' values: $S_0 = 50$, $\mu = 5\%$, $r = 3\%$, $\sigma = 20\%$, $K = 50$, $T = 10$. Number of time steps: 5000.

3.4 Delta in the Black-Scholes Model

The option value evolution depends on the Greeks: the partial derivatives of the option value with respect to its various variables (S_t and T) and parameters[3] (σ and r). The Greeks are a central matter to risk management. Indeed, a Taylor series expansion of C_t shows how the change in the call option value is related to the its Greeks (Hull, 2012):

$$\Delta C_t = \frac{\partial C_t}{\partial S_t}\Delta S_t + \frac{\partial C_t}{\partial T}\Delta T + \frac{\partial C_t}{\partial r}\Delta r + \frac{\partial C_t}{\partial \sigma}\Delta \sigma + \frac{1}{2}\frac{\partial^2 C_t}{\partial S_t^2}(\Delta S_t)^2 + \cdots \tag{3.10}$$

The Greeks of equation (3.10) are, from left to right: Delta, Theta, Rho, Vega and Gamma. In this thesis, however, we will only need to compute the first Greek - Delta - in our empirical analysis to calculate the payoff replication errors.

The Delta of a call option in the BS framework is given by:

$$\Delta_t^{BS} = \mathcal{N}(d_1) > 0 \tag{3.11}$$

As one might expect, the Delta of a call is always positive, and takes value in [0,1]. Also, because it can be shown that Gamma is always positive, we have that the Delta is a strictly increasing function of the stock price. Finally, as expiry approaches, Delta becomes similar to the binary-valued payoff at expiration. Just before expiry, it will be almost zero for S_t just a little below K and almost 1 just above K (Joshi, 2008).

[3]The partial derivative with respect to the strike price K is not a risk factor of an option given that it is set for the whole life of the contract.

3.5 Empirical Limitations

The BS formula, because of its strong assumptions, and especially the constant variance of stock returns, is empirically poorly performing for pricing purposes. Indeed, the stylized facts discussed in section 2.3 are not modeled by the GBM. From Macbeth and Merville (1979), Duan (1995) and Hsieh and Ritchken (2005), we can hold some systematic biases of the BS formula when applied to equities:

- Deep-in-the-money (DITM) options are underpriced.
- At-the-money (ATM) and out-of-the-money (OTM) options are overpriced.
- Options on low-volatility securities are underpriced.
- Short-maturity options are underpriced.

The two first moneyness biases increase with the maturity of the contract. To better understand how these biases can be resolved by more advanced models, we dig into two concepts: the volatility smile (and skew) and the volatility term structure.

To understand what the volatility smile is, we must first precise the concept of implied volatility. From Christoffersen and Jacobs (2003), we define the implied volatility as $\sigma^{imp} = BS^{-1}(C_t^{mkt}, S_t, r, K, T)$, meaning that it is the value of volatility such that the market price of the option is equal to the BS price. As explained by van der Wijst (2013), implied volatilities are commonly used in practice: option traders quote option prices in implied volatilities rather than in raw amount. This measure of volatility is more forward-looking than historical measures.

When we plot the implied volatilities for different levels of the strike price, we would except to get a flat line if the BS formula was correct. However, in practice, with a stock option, what we get is different:

- ITM options tend to have higher implied volatilities than ATM options.
- OTM options tend to have lower implied volatilities than ATM options.

This decreasing relation is called the *volatility skew*. This is another way of specifying the moneyness biases explained before. Figure 3.2 illustrates the volatility skew based on S&P500 call options on three different maturities. We see that the three curves are

FIGURE 3.2: Implied volatility skew of S&P500 index call options.

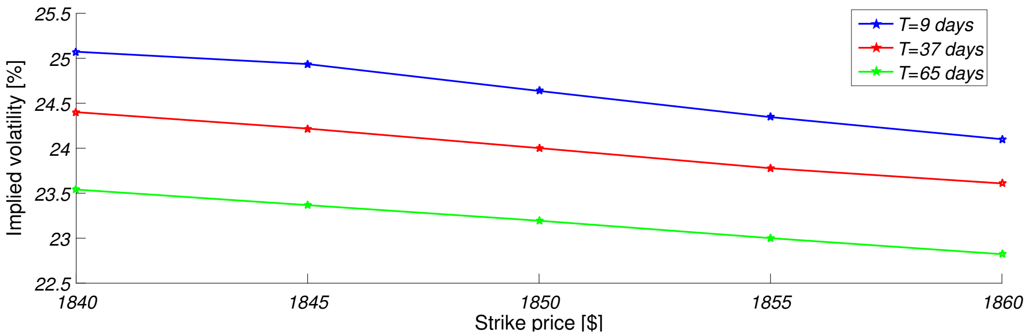

Notes. The initial date is February 10, 2016. The index close price on that date was 1851.86. The implied volatilities are calculated based on the mid bid-ask call option quotes.

decreasing and that the implied volatility increases when the option approaches expiration, which confirms the short-maturity under-pricing mentioned before.

Another common relation that is observed in practice is a U-shaped relation between the implied volatility and the strike price: the well-known *volatility smile*. Hull (2012) explains that this shape is more common for foreign currency options than for equities. However, several authors find evidence of U-shaped relations on stock options (see Macbeth and Merville (1979) or Dennis and Mayhew (2000) for instance).

To understand why the volatility smile effect appears, it is easier to look at the implied (from option quotes) risk-neutral density function of stock returns (see Hull (2012) for technical details). The volatility smile translates into fatter tails than the Gaussian. For stocks, the volatility skew translates into a negatively skewed distribution. These are nothing but the stylized facts explained in section 2.3, that we can now interpret in the context of option pricing. The more advanced models analyzed in this thesis will all try to better fit the asset distribution implied from option prices.

Finally, in addition to the strike price, the implied volatility also depends on the option maturity, which gives rise to the *volatility term structure*. Hull (2012) explains that the relation tends to be positive when short-dated volatilities are historically low and negative when historically high (as in figure 5.2). This is because, in the first case, there is an expectation that volatilities will increase, and the contrary in the second case. If you combine the volatility term structure with the volatility smile/skew, you get a volatility surface that is at the core of the Practitioner Black-Scholes model.

Chapter 4

Alternative Option Pricing Models

4.1 The Practitioner Black-Scholes Model

4.1.1 Model Framework & Calibration

The Practitioner Black-Scholes (PBS) model (also known as ad-hoc BS model) was developed by Dumas et al. (1998). It relies on the considerations of section 3.5 regarding the volatility smile/skew and term structure. The principle is to explicitly model the implied volatility through a deterministic volatility function depending on the strike price and/or the time to maturity. The PBS model is part of the class of *local volatility models*. The most general model investigated by the authors is

$$\sigma^{imp} = \theta_0 + \theta_1 K + \theta_2 K^2 + \theta_3 T + \theta_4 T^2 + \theta_5 KT + \varepsilon \tag{4.1}$$

After estimation of the vector of parameters θ from cross-sections of option prices (see below), the estimate of the option price is obtained as $\hat{C}^{PBS} = C^{BS}\big(\sigma^{imp}(\hat{\theta})\big)$. Using the BS formula in such a way may seem internally inconsistent. However, Berkowitz (2002) showed that the frequent re-calibration of the model "can be viewed as a functional approximation to the true but unknown option pricing formula." Intuitively, market completeness is not verified: for the exact same underlying, two options with two different strike prices/maturities will use a different value for the underlying volatility.

In this thesis, following Dumas et al. (1998), we assume a structural form of the implied volatility that only depends on the strike price. The reason is that the authors found that by including maturity, the PBS model does not perform better in term of in-sample valuation and performs worse in terms of out-of-sample prediction and hedging. However, we assume a functional form which is different from their paper:

$$\sigma^{imp} = \theta_0 + \theta_1 \frac{K}{S_t} + \theta_2 \Big(\frac{K}{S_t}\Big)^2 + \varepsilon \tag{4.2}$$

i.e. we normalize the strike price by the stock price rather than simply using the strike price itself. The rationale is that the volatility smile is in practice often calculated as the relation between the implied volatility and K/S_t because this makes the smile more stable (Hull, 2012). Moreover, we believe that this can bring some flexibility to hedging since the implied volatility will change over the life of the option depending on the value of S_t while it would stay constant if it only depended on K.

Finally, we calibrate the parameter vector θ similarly to the BS model:

$$\hat{\theta} = Arg \min_{\theta} \$\text{MSE}(\theta) \equiv Arg \min_{\theta} \frac{1}{n} \sum_{i=1}^{n} \big(C_i^{mkt} - C_i^{BS}(\sigma_i^{imp}(\theta)\big)^2 \tag{4.3}$$

where n is the number of option contracts in the sample considered. In practice, to ensure that the volatility used in the BS formula is positive, Christoffersen (2012) explains that the structural form used for σ_i^{imp} in equation (4.3) is $\max(0.01, \sigma_i^{imp}(\theta))$.

4.1.2 Empirical Performance

Dumas et al. (1998) found that that their model significantly improves the cross-sectional fit and eliminates an important part of the BS biases. This can be understood by looking at figure 4.1: the implied S&P500 index price density under the PBS model exhibits the stylized facts of equities: negative skewness and fat tails.

Additionally, the PBS model compares favourably to other more advanced classes of models presented thereafter. Yung and Zhang (2003) find that PBS performs better in terms of hedging than the EGARCH model and Christoffersen and Jacobs (2003) show that the PBS model performs somewhat better than the Heston model in-sample and out-of-sample when using consistent evaluation and estimation loss functions. However, a drawback is that Dumas et al. (1998) find that the estimation of the volatility function is unstable through time. This unstability makes PBS less performing in terms of hedging than the BS model in their study. The reason is that the valuation errors of the BS model are stable through time or at least strongly serially dependent. Our empirical study of hedging performance will however disagree with the authors' finding.

FIGURE 4.1: Risk-neutral price densities under the PBS model implied from option quotes on the S&P500 index.

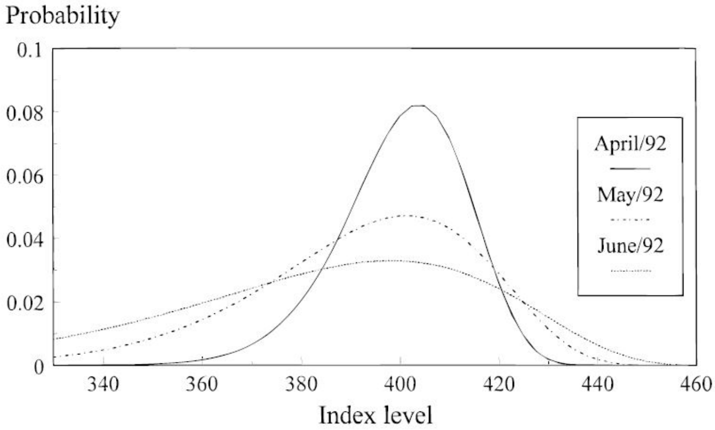

Notes. The figure is taken from Dumas et al. (1998). The initial date was April 1, 1992. The densities are based on parameter estimates of the implied volatility function $\sigma^{imp} = \max(0.01, \theta_0 + \theta_1 K + \theta_2 K^2 + \theta_3 T + \theta_4 T^2 + \theta_5 KT)$.

4.2 The Heston-Nandi GARCH Model

The model of Heston and Nandi (2000) developed in this section is part of the broader class of GARCH option pricing models. Hence, we will first review GARCH modelling and the general GARCH option pricing framework before detailing the HN model.

4.2.1 Modelling Volatility with GARCH

The GARCH (Generalized Auto-Regressive Conditional Heteroskedastic) model was introduced by Bollerslev (1986) who developed the well known GARCH(p,q) model. This model allows for the conditional variance to be a function of lagged conditional variances as well as past sample variances as follows:

$$\begin{aligned} \varepsilon_t | \mathcal{F}_{t-1} &\sim \mathcal{N}(0, h_t) \\ h_t &= \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2 + \sum_{i=1}^{p} \beta_i h_{t-i} \end{aligned} \tag{4.4}$$

where ε_t is a discrete-time random process. We further assume a daily time step. This model is known to incorporate the stylized facts of financial time series: volatility clustering, the mean-reverting behaviour of volatility and the leptokurtic distribution of asset prices. Following the work of Bollerslev, an extensive literature on this class of models has emerged proposing different specifications of the volatility function h_t. A review can be found in Bollerslev et al. (1992). One important technical feature of GARCH models is that for the conditional variance to be stationary (i.e. for the unconditional variance of h_t to be finite), the parameters typically have to be constrained.[1] They also have to be positive for the variance to be positive as well.

4.2.2 GARCH Option Pricing Framework & HN model

Duan (1995) was the first to develop an option pricing model based on the GARCH framework. This model differentiates from the other models of this thesis in that it is a discrete-time model rather than continuous-time. Following his paper, we assume that the daily rate of return under the physical probability measure $\mathbb{P}$ follows the dynamic

$$\ln \frac{S_t}{S_{t-1}} \equiv R_t = r + \delta_t - \frac{1}{2} h_t + \sqrt{h_t} \varepsilon_t \tag{4.5}$$

where $\varepsilon_t | \mathcal{F}_{t-1} \sim \mathcal{N}(0,1)$ and δ_t is the asset risk premium. The volatility process takes the general form

$$h_t = f(h_s, \varepsilon_s; s < t; \theta) \tag{4.6}$$

for some parameter set θ. It is worth noting that GARCH is a *deterministic volatility* framework[2], with the same source of randomness (i.e. ε_t) impacting the return and variance.[3] Stochastic volatility models will be introduced in section 4.3.

The important choice of the model concerns the volatility dynamic function $f(\cdot)$. Duan (1995) used the standard GARCH(p,q) process. However, other authors chose other specifications. Yung and Zhang (2003) analyzed the hedging performance of the EGARCH process of Nelson (1991) that has the appealing feature of removing the parameter restrictions of the traditional GARCH models. Christoffersen and Jacobs (2004), by comparing various GARCH models, found that the out-of-sample pricing analysis favors a rather simple model that, besides volatility clustering, only allows for a standard leverage effect. This model is the NGARCH model. Heston and Nandi (2000) developed a closed-form solution for European options under a GARCH specification. Their model is referred to as the HN model. In this thesis, even though Hsieh and Ritchken (2005)

[1] Bollerslev (1986) for example showed that his model was stationary if $\sum_{i=1}^{q} \alpha_i + \sum_{i=1}^{p} \beta_i < 1$.

[2] Deterministic because the conditional variance $h_t | \mathcal{F}_{t-1}$ is known at time t.

[3] However, as shown by Duan (1995), the GARCH model still allows for negative correlation between the shock component ε_t and the variance process h_t.

showed that the NGARCH model performs better than the HN model, we will adopt the model developed by Heston and Nandi (2000). The rationale behind this choice is that, in practice, a trade-off must be found between the complexity of the model and the ease of implementation (i.e. the computational cost of the model). The analytical tractability of the HN model through its closed-form formula makes it very attractive from a computational perspective: we do not have to resort to Monte Carlo simulation to calibrate the parameters to market quotes.

In the HN model, the risk premium is defined as $\delta_t = (\lambda + \frac{1}{2})h_t$. This means that the return process in equation (4.5) becomes

$$R_t = r + \lambda h_t + \sqrt{h_t}\varepsilon_t \tag{4.7}$$

with the volatility dynamic being given by

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2\big(\varepsilon_{t-1} - \gamma\sqrt{h_{t-1}}\big)^2 \tag{4.8}$$

with $\beta_0, \beta_1, \beta_2 > 0$. The unconditional variance $\mathbb{E}^{\mathbb{P}}[h_t]$ is finite and equals $\frac{\beta_0+\beta_2}{1-\beta_1-\beta_2\gamma^2}$ if $\beta_1 + \beta_2\gamma^2 < 1$ (in which case the model is stationary). The left-hand side of this inequality is called the persistence of model. It is a measure of how fast the process h_t reverts to the unconditional variance. The closer it is to 1, the lower it will revert and the higher will be the volatility clustering effect. The parameter γ determines the skewness of the asset return distribution, a positive value indicating a negative skewness. The heteroskedasticity finally implies that the unconditional returns distribution is leptokurtic. Figure 4.2 shows a simulation of the HN stock price (under the physical measure), compared to a GBM process. The gap between the two paths increases with time, but HN does not appear to be more volatile than GBM, possibly because both models are affected by only one source of random shocks. Figure 4.3 displays the density function of stock returns under the HN model, compared to the normal density. The density under the HN model depicts the stylized facts of stock returns: negative skewness (-0.38), a fatter left tail and a less fat right tail than the normal distribution.

To develop the GARCH option pricing model, Duan (1995) generalized the conventional risk-neutral valuation relationship. He provided a locally risk-neutral valuation relationship (LRNVR) that holds under some familiar assumptions on preferences and that is

FIGURE 4.2: Comparison of stock price simulations under the GBM and HN processes.

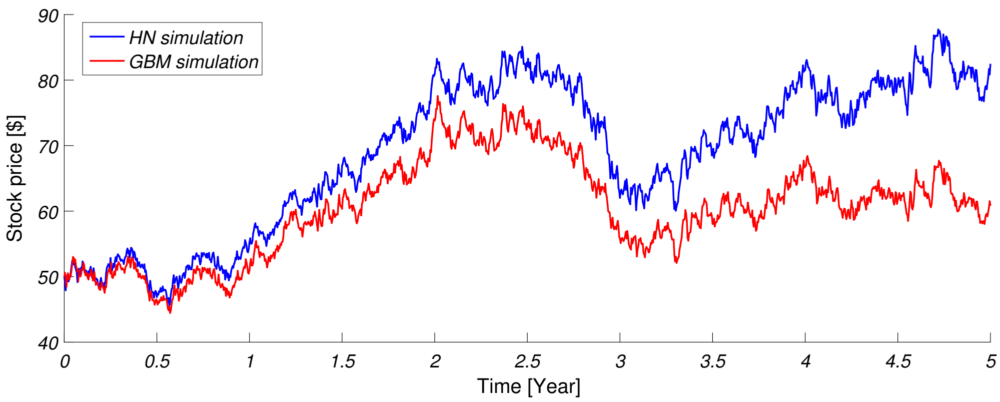

Notes. The simulations were done under the physical measure, and the GBM process was simulated with a Milstein discretization scheme. The same Gaussian random shocks were applied for each model. The GBM process instantaneous (annualized) mean and volatility are reflected by the HN parameters' values as follows: λ is set such that $\mu = 252 * (r + \lambda\bar{h})$, and β_2 such that $\sigma = \sqrt{252\bar{h}}$ where $\bar{h} = \beta_0/(1 - \beta_1 - \beta_2\gamma^2)$, i.e. the long run variance under the HN model. The factor 252 reflects the daily time step used to simulate the HN process. The parameters' values are: $S_0 = 50$, $h_0 = \bar{h} = 1.59 * 10^{-4}$, $\mu = 5\%$, $r = 0.03/252$, $\sigma = 20\%$, $\beta_0 = 8 * 10^{-6}$, $\beta_1 = 0.45$, $\beta_2 = 1 * 10^{-6}$, $\gamma = 700$, $\lambda = 0.5$, $T = 5$.

FIGURE 4.3: Comparison of the physical density of stock returns under the HN model with the normal distribution.

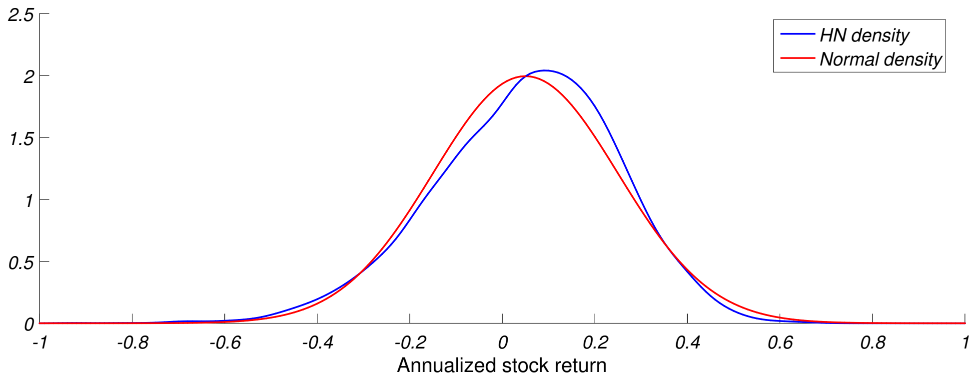

Notes. The HN density was obtained by Monte Carlo simulation, with a Gaussian kernel under Matlab default bandwidth. The normal distribution's mean and standard deviation are reflected by the HN parameters' values as follows: λ is set such that $\mu = 252 * (r + \lambda\bar{h})$, and β_2 such that $\sigma = \sqrt{252\bar{h}}$ where $\bar{h} = \beta_0/(1 - \beta_1 - \beta_2\gamma^2)$, i.e. the long run variance under the HN model. The factor 252 reflects the daily time step used to simulate the HN process. The parameters' values are: $S_0 = 50$, $h_0 = \bar{h} = 1.59 * 10^{-4}$, $\mu = 5\%$, $r = 0.03/252$, $\sigma = 20\%$, $\beta_0 = 8 * 10^{-6}$, $\beta_1 = 0.45$, $\beta_2 = 1 * 10^{-6}$, $\gamma = 700$, $\lambda = 0.5$, $T = 1$. Number of Monte Carlo random paths: 5000.

satisfied by a measure $\mathbb{Q}$ if

$$\mathbb{E}^{\mathbb{Q}}[\exp(R_t)|\mathcal{F}_{t-1}] = \exp(r) \tag{4.9}$$

meaning that the expected return is the risk-free rate, and

$$\mathrm{Var}^{\mathbb{Q}}[R_t|\mathcal{F}_{t-1}] = \mathrm{Var}^{\mathbb{P}}[R_t|\mathcal{F}_{t-1}] = h_t \tag{4.10}$$

which means that the conditional variances under the two measures are equal. The LRNVR yields a well-specified model that does not locally depend on preferences. However, it is insufficient for fully eliminating the preference parameters, but strong enough to reduce them to the risk premium parameter λ.

Under the risk-neutral measure $\mathbb{Q}$, the return process evolves as

$$R_t = r - \frac{1}{2}h_t + \sqrt{h_t}\varepsilon_t^* \tag{4.11}$$

with $\varepsilon_t^*|\mathcal{F}_{t-1} \sim \mathcal{N}(0,1)$. The volatility process becomes

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2\big(\varepsilon_{t-1}^* - \omega\sqrt{h_{t-1}}\big)^2 \tag{4.12}$$

with $\omega = \gamma + \lambda + \frac{1}{2}$. The risk-neutral pricing measure is determined by four parameters that have to be estimated: $\beta_0, \beta_1, \beta_2$ and ω. Additionally, following the procedure of Hsieh and Ritchken (2005), we take h_0 as a parameter to be optimally determined from the data since it is not directly observable from the market. Next section explains the implementation of the HN model to calibrate the parameters based on market quotes.

4.2.3 Implementation of the HN model

In their paper, Heston and Nandi (2000) derived a closed-form formula for the option price. This is a very attractive feature since it makes the implementation much less cumbersome than a Monte Carlo approach (Joshi, 2008). Following the reasoning of Heston and Nandi (2000), we write the value of a European call option as follows:

$$C_t^{HN} = S_t P_1 - Ke^{-r(T-t)}P_2 \tag{4.13}$$

The two terms P_1 and P_2, as explained by Gatheral (2006), have the same meaning as in the BS model (see section 3.2.1). Rouah and Vainberg (2007) explain that P_1 and P_2 can be computed by integrating the characteristic function of the log price $f(i\phi)$:[4]

$$
\begin{aligned}
P_1 &= \frac{1}{2} + \frac{e^{-r(T-t)}}{\pi S_t} \int_0^{\infty} \mathbb{R}\left[\frac{Ke^{-i\phi} f(i\phi+1)}{i\phi}\right] d\phi \\
P_2 &= \frac{1}{2} + \frac{1}{\pi} \int_0^{\infty} \mathbb{R}\left[\frac{Ke^{-i\phi} f(i\phi)}{i\phi}\right] d\phi
\end{aligned}
\tag{4.14}
$$

where i is the imaginary unit and $\mathbb{R}[\cdot]$ refers to the real part of a complex function.[5]

The characteristic function takes the following form:

$$
f(i\phi) = S_t^{i\phi} \exp(A_t + B_t h_t) \tag{4.15}
$$

where

$$
\begin{aligned}
A_t &= A_{t+1} + \phi r + B_{t+1}\beta_0 - \frac{1}{2}\ln(1 - 2\beta_2 B_{t+1}) \\
B_t &= \phi\Big(\omega - \frac{1}{2}\Big) - \frac{1}{2}\omega^2 + \beta_1 B_{t+1} + \frac{\frac{1}{2}(\phi-\omega)^2}{1 - 2\beta_2 B_{t+1}}
\end{aligned}
\tag{4.16}
$$

Both equations can be solved recursively from time T by using the terminal conditions $A_T = B_T = 0$. From the option price formula, we calibrate the parameters by minimizing the pricing \$MSE, similarly to Heston and Nandi (2000).

4.2.4 HN Delta & Market Completeness

From equation (4.13), we easily get that the Delta Δ_t^{HN} for a call option is

$$
\Delta_t^{HN} = P_1 > 0 \tag{4.17}
$$

Duan (1995) explains that GARCH adds another level of flexibility to dynamic hedging, the option position being smaller (higher) in a low (high) variance state than with BS.

It is worth noting that market completeness is not verified by GARCH models: perfect replication of every payoff via a self-financing strategy is not possible. The reason has to do with "the poor cohabitation between discrete time modeling and the infinite number of states of the innovations"[6] (Ortega, 2009). This means that the risk-neutral measure used here is not unique.

[4]The characteristic function $f(\phi)$ is the moment generating function of the log price under the risk-neutral measure: $f(\phi) = \mathbb{E}_t^{\mathbb{Q}}[S_T^{\phi}]$

[5]To improve the efficiency of the call option price calculation with Matlab, we follow the approach of Rouah and Vainberg (2007) who combine the two integrals into a single one.

[6]The innovations refer to the random shocks ε_t.

4.2.5 Empirical Performance

Hsieh and Ritchken (2005) showed that, in-sample, the HN model removes a significant portion of the biases for each moneyness-maturity category. Out-of-sample, they find that the pricing errors remain reasonable, except for DOTM contracts where "the HN model is unable to explain the volatility strike bias." Heston and Nandi (2000) showed that the HN model outperforms the PBS model both in terms of in-sample fitting and out-of-sample forecasting, a finding that we will confirm in the empirical part.

Regarding hedging performance, Yung and Zhang (2003) showed that the EGARCH model performs worse than the PBS model, irrespective of moneyness categories and hedging horizons. The authors explain that one plausible reason is that "more parameters in a model may result in more freedom in fitting option price curvature, but may generate less accurate Delta estimates." Our empirical study will strongly confirm this statement. Finally, Petitjean and Moyaert (2011) found that HN led to lower hedging errors than Heston during the volatile subprime crisis. However, our own findings will point otherwise, though less strikingly under a higher underlying volatility.

Next section aims at analyzing the class of so-called stochastic volatility models that are in many ways theoretically related to GARCH models.

4.3 The Heston Stochastic Volatility Model

4.3.1 Theoretical Framework & Heston Model

Stochastic volatility (SV) models consider the volatility of the underlying asset as being a continuous-time stochastic diffusion process. The process contains a second source of risk (dW_V) affecting the level of instantaneous volatility (Jackel, 2004). Following Joshi (2008), the general form of these models is

$$\begin{aligned} dS_t &= \mu S_t dt + \sqrt{V_t} S_t dW_S \\ dV_t &= \mu_V dt + \sigma_V V_t^{\alpha} dW_V \end{aligned} \tag{4.18}$$

with α a positive real number. The two Brownian motions dW_S and dW_V may be correlated or uncorrelated. The risk-neutral version of the model is

$$dS_t = rS_t dt + \sqrt{V_t} S_t dW_S$$
$$dV_t = \tilde{\mu}_V(S_t, V_t, t)dt + \sigma_V V_t^{\alpha} dW_V \tag{4.19}$$

where $\tilde{\mu}_V$ is an arbitrary function of (S_t, V_t, t). Indeed, as explained by Joshi (2008): "the drift of the variance process is not relevant to being a martingale, as its size will only magnify up and down moves, not change the mean of the stock value." This reflects an important feature of SV model: volatility is not a directly tradable quantity. Therefore, the market has two different sources of uncertainty but only one underlying and so, as it was the case with GARCH models, SV models imply an incomplete market. More precisely, the replication strategy developed in section 3.3 that relies only on the stock and a cash account is not sufficient to perfectly hedge the option: another derivative is needed to hedge against the volatility risk.[7]

Many different specifications for the volatility have been developed in the literature.

Hull and White (1987) modeled the variance V_t with a log-normal process of a form similar to the asset price process, with two correlated Brownian motions. However, this model has some drawbacks. Solving the model requires complex numerical methods (Kermiche, 2014) and the dynamics of the model predicts that "both expectation and most likely value of instantaneous volatility converge to zero" (Jackel, 2004).

Stein and Stein (1991) developed a model where a mean reversion process is used for the volatility. However, one drawback is their assumption that volatility is not correlated with the underlying, which does not allow for skewness (Kermiche, 2014). Also, the model implies that the level of volatility has its most likely value at zero (Jackel, 2004).

In this thesis, we follow the famous model developed by Heston (1993) who relaxed the model of Stein and Stein (1991) by allowing the volatility and asset price to be correlated. The volatility diffusion process under the risk-neutral measure follows a Cox-Ingersoll-Ross process (Cox and Ross, 1985) that allows for mean-reversion:

$$dV_t = \kappa(\theta - V_t)dt + \alpha\sqrt{V_t} dW_V \tag{4.20}$$

[7] See Liu and Pan (2003) for a market complete framework under the Heston model where the investment universe is composed of a derivative security in addition to the underlying asset and the bond process.

with $E[dW_S \cdot dW_V] = \rho dt$.[8] If we denote κ^* and θ^* to be the corresponding drift parameters of the process under the physical measure, we have the following relations:

$$\begin{aligned} \kappa &= \kappa^* + \lambda \\ \theta &= \frac{\kappa^* \theta^*}{\kappa^* + \lambda} \end{aligned} \tag{4.21}$$

where λ is the volatility risk premium. These relations are important when going from one measure to another. The process for V_t means that the instantaneous variance reverts to the long-run variance θ at a rate given by κ. This model allows for skewness and kurtosis in the distribution of the underlying. Bakshi et al. (1997) indicate that, in SV models, the correlation between the volatility and the underlying (ρ) determines the level of skewness, while the volatility of volatility (α) determines the level of kurtosis. Figure 4.4 displays a simulated stock price under the Heston model compared to a GBM process. The Heston model produces a much more volatile stock price process, especially compared to what we observed for HN in figure 4.2, possibly because Heston is affected by two sources of random shocks. Figure 4.5 shows the density function of stock returns under the Heston model, compared to the normal density. We find similar features than for HN: the density depicts negative skewness (-0.90), a fatter left tail and a less fat right tail than the normal. However, it appears that Heston is much more flexible than HN in capturing a large negative skewness.

This last observation is confirmed by Kermiche (2014) who asserts that the attractiveness of the Heston model is its ability to calibrate to the market given skew. However, Jackel (2004) explains that this requires an unrealistic negative correlation, with ρ almost always between -0.7 and -1. This unrealistic range is also necessary to fairly effectively price short-term options. A very low value of κ is also required.

A final implication is that the volatility can become very low or even negative due to the randomness of its process. Specifically, it can become negative unless $2\kappa\theta > \alpha^2$. This is the Feller condition (Feller, 1951).[10] Due to the square root process of the variance,

[8] From Hull (2012), we can generate two correlated Brownian motions by setting $dW_S = \rho dW_1 + \sqrt{1-\rho^2} dW_2$ and $dW_V = dW_1$ where dW_1 and dW_2 are two uncorrelated Brownian motions.

[9] From Rouah (n.a.), we have that the Milstein discretization for the Heston model is, under the physical measure: $S_{t+\Delta t} = S_t + \mu S_t \Delta t + \sqrt{V_t \Delta t} S_t Z_S + \frac{1}{4} S_t^2 \Delta t (Z_S^2 - 1)$ and $V_{t+\Delta t} = V_t + \kappa^*(\theta^* - V_t)\Delta t + \alpha \sqrt{V_t \Delta t} Z_V + \frac{1}{4}\alpha^2 \Delta t (Z_V^2 - 1)$ where $Z \sim \mathcal{N}(0,1)$ and Δt is the discrete time step.

[10] When simulating the Heston model under a Milstein discretization scheme, as in figures 4.4 and 4.5, then as shown by Gatheral (2006), the inequality becomes less constraining: $4\kappa\theta > \alpha^2$.

FIGURE 4.4: Comparison of stock price simulations under the GBM and Heston processes.

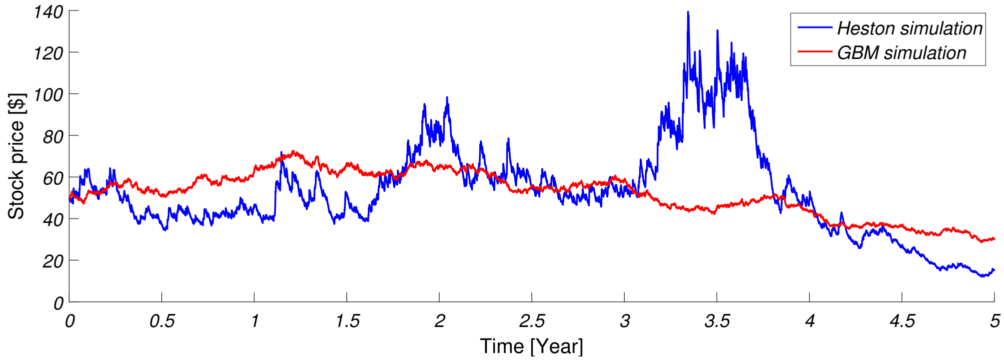

Notes. The simulations were done under the physical measure with a Milstein discretization scheme.[9] The GBM process instantaneous volatility equals the long-run volatility of the Heston model. The same Brownian shocks were applied to the GBM process and the stock price process of the Heston model. Parameters' values: $S_0 = 50$, $V_0 = 4\%$, $\mu = 5\%$, $\sigma = 20\%$, $\theta^* = 4\%$, $\kappa^* = 0.3$, $\alpha = 0.2$, $\rho = -0.8$, $T = 5$. Number of time steps: 5000.

FIGURE 4.5: Comparison of the physical density of stock returns under the Heston model with the normal distribution.

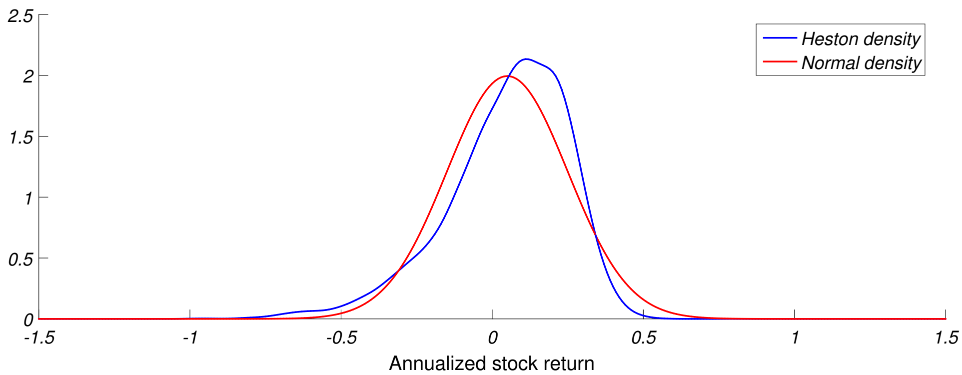

Notes. The Heston density was obtained by Monte Carlo simulation, with a Gaussian kernel under Matlab default bandwidth. The simulation was done with a Milstein discretization scheme. The normal distribution's standard deviation equals the long-run volatility of the Heston model. Parameters' values: $V_0 = 4\%$, $\mu = 5\%$, $\sigma = 20\%$, $\theta^* = 4\%$, $\kappa^* = 0.3$, $\alpha = 0.2$, $\rho = -0.8$, $T = 1$. Number of time steps: 500. Number of Monte Carlo random paths: 5000.

this inequality will be enforced when solving the model. Jackel (2004) explains that this possibility of a very low variance, combined with the low value of κ, implies that "the volatility can stay extremely low or very high for long periods of time."

4.3.2 Implementation of the Heston Model

Heston (1993) derived a closed-form solution for the option price under his model. Similarly to the HN model, we first write the value of a European call option as follows:

$$C_t^{SV} = S_t P_1 - K e^{-r(T-t)} P_2 \tag{4.22}$$

Then, we follow the implementation approach of Crisostomo (2014) who provides efficient and simple solutions to calibrate the Heston model. In his paper, he shows that P_1 and P_2 can be found by solving the following inverse Fourier transforms:

$$\begin{aligned} P_1 &= \frac{1}{2} + \frac{1}{\pi}\int_0^{\infty} \mathbb{R}\left[\frac{e^{-iw\ln(K)}\Psi_{\ln S_T}(w-i)}{iw\Psi_{\ln S_T}(-i)}\right]dw \\ P_2 &= \frac{1}{2} + \frac{1}{\pi}\int_0^{\infty} \mathbb{R}\left[\frac{e^{-iw\ln(K)}\Psi_{\ln S_T}(w)}{iw}\right]dw \end{aligned} \tag{4.23}$$

where $\Psi_{\ln S_T}(w)$ is the characteristic function of the log price.[11] It is interesting to note that this approach is not specific to the Heston model: it can be applied for any underlying process S_t whose characteristic function is known. For the Heston model, the characteristic function is derived by Crisostomo (2014) and is given by

$$\Psi_{\ln S_t}(w) = \exp\big(C(t,w)\theta + D(t,w)V_0 + iw\ln(S_0 e^{rt})\big) \tag{4.24}$$

with

$$\begin{aligned} C(t,w) &= \kappa\left[r^- t - \frac{2}{\alpha^2}\ln\left(\frac{1-ge^{-ht}}{1-g}\right)\right] \\ D(t,w) &= r^- \frac{1-e^{-ht}}{1-ge^{-ht}} \\ r^{\pm} &= \frac{\beta \pm h}{\alpha^2}; \quad h = \sqrt{\beta^2 - 4\eta\gamma} \\ g &= \frac{r^-}{r^+}; \quad \eta = -\frac{w^2}{2} - \frac{iw}{2}; \quad \beta = \kappa - \rho\alpha i w; \quad \gamma = \alpha^2/2 \end{aligned} \tag{4.25}$$

The calibration to option quotes will finally be done with a \$MSE loss function. As for the HN model, V_0 is optimally determined in addition to all the other parameters.

4.3.3 Heston Delta

From Bakshi et al. (1997), we can find the closed-form formulas of the Greeks of the model, adapted to the implementation approach of Crisostomo (2014) presented in the previous section. Similarly to HN, the Delta of a call option is given by

$$\Delta_t^{Heston} = P_1 > 0 \tag{4.26}$$

[11] In a continuous setting, the characteristic function of a stochastic process X is the Fourier transform of its probability density function: $\Psi_X(w) = \mathbb{E}[e^{iwX}] = \int_{-\infty}^{+\infty} e^{iwx} f(x)dx$

4.3.4 Link with GARCH Models

SV and GARCH models are closely related. Of course, these two frameworks have important differences, mainly that SV is a continuous-time model (vs discrete-time for GARCH) with a stochastic volatility that acts as a second source of noise (vs deterministic volatility and only one source of noise for GARCH). The previous simulations of each model also showed that Heston is much more volatile. However, both models specify a volatility function that is time-varying and asymmetric (Lehar et al., 2002). In fact, Heston and Nandi (2000) showed that their model converges to Heston as the length of the time interval tends to zero. Chapter 6 will provide empirical insights into how differently or closely these two models behave and perform regarding payoff replication. We will show that, though they perform very closely, HN is not a good alternative to Heston for hedging purposes.

4.3.5 Empirical Performance

Bakshi et al. (1997) showed that "taking stochastic volatility into account is of the first-order importance in improving upon the BS formula." They also found out that, for hedging purposes, adding more stochastic components such as a jump process to the SV model does not improve its performance, and that the only real improvement brought by SV models over BS emerges for OTM calls, which the results on Apple will confirm. Petitjean and Moyaert (2011) found that the Heston model outperforms the BS model in terms of hedging, even though, as mentioned in section 4.2.5, it leads to larger hedging errors than the HN model. Our findings will contradict both of these findings. In terms of pricing fit, they also conclude that the Heston model seems to perform better for medium or long-term options. For short-term options, they find that the Bates model (which includes jumps) works best. Even though we will not analyze jump models in this thesis, it is important to note that, as shown by Bakshi et al. (1997), "adding the random jump feature to the SV model can further improve its performance, especially in pricing short-term options." Finally, Christoffersen and Jacobs (2003) found that Heston is slightly outperformed by PBS for pricing when using consistent estimation and evaluation loss functions. Conversely, we will show that Heston quite often outperforms PBS for out-of-sample pricing, except at times for long-term or far-from-the-money options.

Chapter 5

Methodology

This chapter details the methodology that we follow to conduct our studies of the payoff replication performance of the four models previously presented, to ultimately provide answers to our research questions presented in chapter 1. First, we describe the data that we collected and how we filtered them. Second, we explain how Delta hedging is measured in the literature and how our payoff replication criterion differs thereof. Third, we look at how we proceeded to do our empirical application by explaining the market calibration procedure, the sliding window technique and the error measures used to classify the models. Finally, we explain how we investigated the over-hedging effect based on simulated data from the HN model. Most of our studies are performed with Matlab. The main functions and scripts that we used are described in appendix A, and can be found on the CD attached to the printing version of the thesis.

5.1 Description of Data

In this thesis, we work with two distinct time series, both corresponding to a different level of volatility. This will enable to verify whether and how our conclusions depend on the level of volatility, keeping in mind that the two series do not only distinguish by the volatility as section 6.1 will show. The call options are of European-style. The two underlying assets' samples considered are the following:

1. *Low volatility sample*: S&P500 index options from January 4 to March 30, 2016. The furthest maturity considered for hedging is May 31, 2016.[1] The annualized standard deviation of log-returns from January 4 to May 31, 2016 is 15.71%.

[1] Section 5.3.3 will explain that some days will be used for parameters estimation purposes and others to calculate the payoff replication errors out-of-sample. Longer maturities will be used for estimation than for payoff replication.

2. *High volatility sample*: Apple stock options from January 4 to March 30, 2016. The furthest maturity considered for hedging is June 17, 2016. The annualized standard deviation of log-returns from January 4 to June 17, 2016 is 27.49%.

We wish to underline that, in the literature, a longer time span is generally analyzed, e.g. 3 years in Christoffersen and Jacobs (2004) and Heston and Nandi (2000), 4 years in Bakshi et al. (1997) or 5 years and a half in Dumas et al. (1998), with data collected every week. However, our option data were obtained on Bloomberg where we only had access to 3 months of historical quotes for options. We therefore collected data on a daily basis rather than weekly to have a large enough sample of options.

In general, the reader should keep in mind that the results in chapter 6 may not always be robust to different underlying asset samples. That said, comparing the results with the literature will enable to confirm previous findings and point to counter-examples. In general, we will get profound insights into the hedging performance and behaviour of our models to come up with relevant implications and suggestions for future research.

Figure 5.1 displays the time series of stock prices and log-returns for the two underlying from January 4, 2016 to the furthest maturity considered for payoff replication. We clearly see that Apple log-returns are more volatile than the S&P500 ones. Apple also had an overall negative return over the period considered, in the contrary to the S&P500. Section 6.1 will give a detailed statistical description of the two time series.

The S&P500 index is a market value weighted index based on 500 highly traded top American companies in leading industries. It is denoted by the ticker SPX (S&P Dow Jones Indices, 2016a). SPX options are traded on the Chicago Board Options Exchange (CBOE). They have a large notional size with a $100 multiplier and are very actively traded with an average daily volume of 938,620 in 2015. They are finally extensively used in the literature (e.g. Dumas et al. (1998) or Christoffersen and Jacobs (2004)).

Apple stock is traded on the NASDAQ stock exchange and is denoted by the ticker AAPL. Options on Apple are traded on the CBOE and are also actively traded.

For each of the two samples, we collected four different types of data:

FIGURE 5.1: Time series of stock prices and log-returns for the S&P500 index and Apple stock.

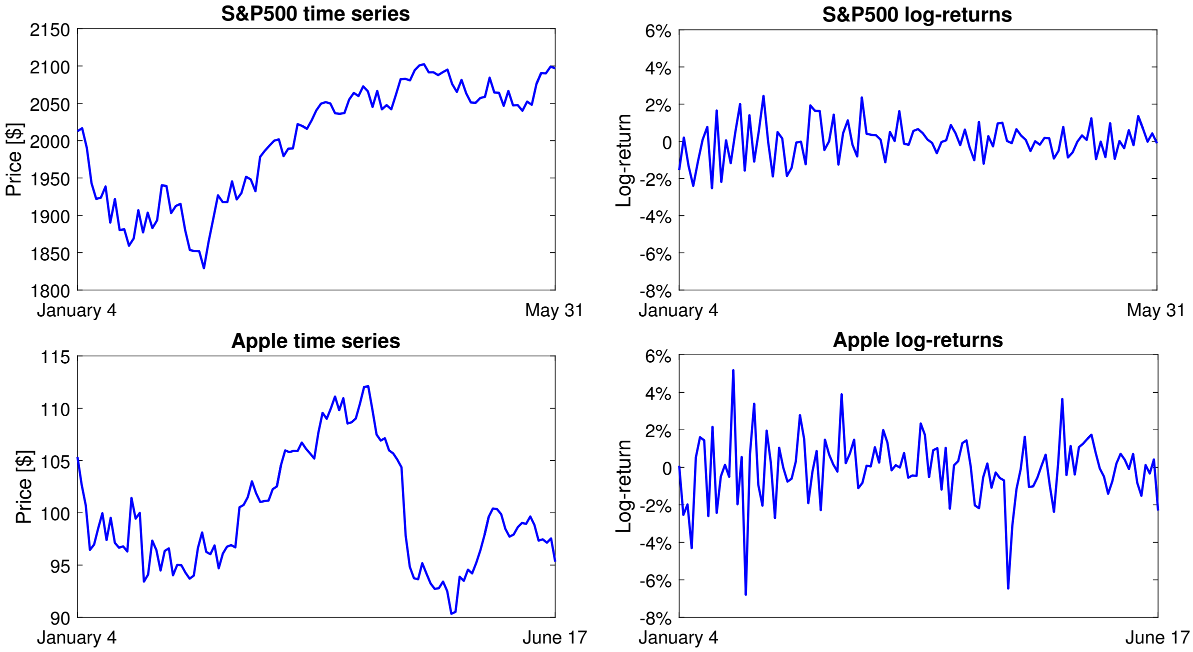

Notes: The time span goes from January 4, 2016, which is the first date of our sample of options, to the furthest maturity considered for payoff replication, i.e. May 31, 2016 for the S&P500 and June 17, 2016 for Apple.

1. *Option quotes.* We collected option data on every day of our sample based on mid bid-ask closing quotes. The quotes were collected on 10 strikes and 5 maturities. This allows to have a rather large number of simultaneous observations. Bates (2003) points out that this allows to compensate for measurement errors in the option data such as synchronization errors with the underlying asset prices[2] or rounding errors. Ten strikes are chosen so as to cover a reasonably wide range of moneyness. For the SPX options, we had access on Bloomberg to the SPX End-of-Month options that expire on the last trading day of each month (CBOE, 2016). It must be noted that wider values of the strike are available on Bloomberg for options with a longer maturity.

2. *Underlying assets.* The data on the S&P500 and Apple time series were obtained on Yahoo Finance based on the closing prices of each day. Because of the methodology by which the S&P500 index is constructed as detailed by S&P Global (2016), the

[2]Non-synchronization arises when the option quote and the corresponding underlying price are not recorded at the same moment.

S&P500 closing prices are adjusted to corporate actions (e.g. stock splits, dividends), while it is not the case for Apple.

3. *Dividends.* The asset prices have to be adjusted to the dividends that accrue during the life of the option. We follow the procedure of Bakshi et al. (1997): the present discounted value of the discrete daily dividends during the life of the option is subtracted from the underlying price. We have access to the S&P500 Dividend Points Index (Annual) on S&P Dow Jones Indices (2016b), which tracks the total daily dividends from the constituents of the S&P500.[3] They are displayed on figure B.1 in appendix. For Apple, the dividends were available on Yahoo Finance. Apple issues dividends every 3 months, so that only two dividends had to be accounted for: 0.52 on February 4, 2016 and 0.57 on May 5, 2016. The present discounted value of the daily dividends is

$$D_{pv} = \sum_{i=1}^{n} D_i e^{-r_i t_i} \tag{5.1}$$

where D_i is the ith cash dividend, r_i is the risk-free rate for the t_i period, t_i is the time to the ith dividend from the initial date, and n is the number of dividend payments during the life of the option.

4. *Interest rate.* Finally, as a proxy to the risk-free interest rate, we use the daily Treasury-bill (T-bill) rates obtained on Board of Governors of the Federal Reserve System (2016). A T-bill is a"short-term debt obligation backed by the U.S. government with a maturity of less than one year" (Investopedia, n.a.). It is commonly used in the literature as a proxy for the short term risk-free rate. To obtain the rate for a particular maturity, we apply the method of Bakshi et al. (1997): we use two T-bill rates that straddle the maturity date to obtain the risk-free rate for that date, by linear interpolation. The T-bill rates we use are the 1-month, 3-month and 6-month rates. They are displayed on figure B.2 in appendix. We transform them into continuously compounded rates.[4]

[3] At the time of the analysis, dividends were not available for some long-term maturities that extended to June 2016 and beyond and that we needed for our estimation windows (see section 5.3.3 for an explanation of the sliding window technique). In that case, we approximated the future dividends with the average dividend in 2015, adjusted to the growth rate of dividends between 2015 and 2016.

[4] The continuously compounded interest rate r^c is calculated as $r^c = \ln(1 + r)$ where r is the discrete interest rate (Chance, 2003).

Based on these data, we apply the following elimination criteria:

- Following Bakshi et al. (1997), options with a price below \$3/8 are removed "to mitigate the impact of price discreteness on option valuation." Indeed, for options with low prices, pricing errors due to rounding may become substantial.
- Following Heston and Nandi (2000) and Hsieh and Ritchken (2005), we exclude options with maturities fewer than six days. Short-term options, as explained by Dumas et al. (1998), have relatively small time premiums and are thus sensitive to non-synchronicity and other possible measurement errors.
- Following Yung and Zhang (2003), we eliminate options that violate the lower bound arbitrage condition, i.e. when $C_t < S_t - D_{pv} - Ke^{-r(T-t)}$ where r is the T-bill rate with maturity date T.
- Finally, as explained in details in section 5.3.3, some days will be used for parameters' estimation purposes, and others to calculate the payoff replication errors out-of-sample. For the latter days, we need the data on all the stock prices until maturity of the option to calculate the errors. Therefore, for those days, we eliminated options with maturity dates extending beyond the time of the analysis. The furthest maturities used for payoff replication are therefore May 31, 2016 for the S&P500 and June 17, 2016 for Apple. However, we did not apply this elimination criterion for the days used for estimation in order to keep a large enough sample size and to be consistent across windows in the range of maturities used in the calibration.

Table C.1 in appendix reports the number of options eliminated for each time series. The total sample refers to the sample before elimination of the options having too distant maturities to calculate the hedging errors[5], and the hedging sub-sample is calculated after eliminating them.

Then, we classify options depending on moneyness (at inception) and days to maturity.

Defining moneyness as $\mathrm{M} = \frac{S_t - D_{pv}}{K}$, we distinguish four moneyness categories that differ for the S&P500 and Apple because larger moneyness levels are quoted for Apple:

- DOTM: M<0.98 for S&P500 (15% of the total sample), M<0.94 for Apple (18%).

[5] We report this number because these long-maturity options will still be used when analyzing the out-of-sample pricing performance in section 6.3.2.

- OTM: 0.98<M<1 for S&P500 (44%), 0.94<M<1 for Apple (22%).
- ITM: 1<M<1.02 for S&P500 (31%), 1<M<1.06 for Apple (26%).
- DITM: M>1.02 for S&P500 (10%), M>1.06 for Apple (34%).

Three levels of maturity are then distinguished:

- Short-term contracts have maturities between 6 and 30 days included, accounting for 26% of both of our options samples.
- Mid-term contracts have maturities between 31 and 60 days included (30%).
- Long-term contracts have maturities of more than 60 days (44%).

Figure 5.1 shows the number of contracts used for each category from the total sample of options for the two time series. We can particularly notice that DOTM and DITM SPX options are not available for short-term contracts, simply because Bloomberg quotes a larger range of strike prices only for options with enough days to maturity. Figure 5.2 shows the average implied volatility in each moneyness-maturity category, clearly showing the BS biases. The volatility smile is especially visible for short-term options (and mid-term ones for the S&P500), while long-term contracts rather feature a volatility skew (and mid-term ones for Apple). Finally, table B.1 in appendix gives more details on the quantiles of maturity, moneyness and call option price.

5.2 Delta Hedging Performance in the Literature

Before explaining the methodology followed to assess the payoff replication performance of our models, we wish to do a brief detour into how Delta hedging performance is measured in the literature. Indeed, as already argued in chapter 1, our hedging criterion differs from the one used in the literature, which is why it is worth briefly explaining how it does differ and why we opted for payoff replication in our thesis.

Typically, hedging errors in the literature are calculated out-of-sample in three steps (see Bakshi et al. (1997) for example). First, a hedged portfolio is formed on date t by shorting a call option, going long in Δ_t shares of the stock, and investing the remaining (denoted X_0) at the risk-free rate (or borrowing if necessary). Second, this portfolio is discretely re-balanced until a future time $t + \Delta t$, where Δt is typically one day. The

TABLE 5.1: Number of option contracts used across moneyness and maturity.

S&P500

	DTM⩽30	30<DTM⩽60	DTM>60	*Total*
M<0.98	0	45	372	417
0.98<M<1	405	502	341	1248
1<M<1.02	315	288	276	879
M>1.02	0	25	261	286
Total	720	860	1250	2830

Apple

	DTM⩽30	30<DTM⩽60	DTM>60	*Total*
M<0.94	35	172	289	496
0.94<M<1	257	141	203	601
1<M<1.06	273	164	248	685
M>1.06	122	330	450	902
Total	687	807	1190	2684

Notes: The two samples of option contracts cover the period January 4, 2016 to March 30, 2016. M refers to moneyness, defined as $\frac{S_t - D_{pv}}{K}$, and DTM to Days To Maturity. We report the statistics on the number of contracts for each moneyness-maturity category for the total sample of options used for each time series. The total sample also includes those options that will have to be eliminated to calculate the payoff replication errors because they have maturity dates extending beyond the time of the analysis.

TABLE 5.2: Average implied volatility across moneyness and maturity.

S&P500

	DTM⩽30	30<DTM⩽60	DTM>60
M<0.98	/	18.074	17.22
0.98<M<1	17.42	17.6	17.21
1<M<1.02	18.51	18.68	18.76
M>1.02	/	22.89	21.56

Apple

	DTM⩽30	30<DTM⩽60	DTM>60
M<0.94	28.84	26.89	27.13
0.94<M<1	27.15	27.80	28.54
1<M<1.06	27.41	29.02	29.32
M>1.06	32.67	31.58	31.05

Notes: The two samples of option contracts cover the period January 4, 2016 to March 30, 2016. M refers to moneyness, defined as $\frac{S_t - D_{pv}}{K}$, and DTM to Days To Maturity. We report the statistics on the average implied volatility (in percent) for each moneyness-maturity category for the total sample of options used for each time series. The total sample also includes those options that will have to be eliminated to calculate the payoff replication errors because they have maturity dates extending beyond the time of the analysis.

hedging error after that time step is calculated as

$$\varepsilon_{t+\Delta t} = \Delta_t S_{t+\Delta t} + X_0 e^{r\Delta t} - C^{mkt}_{t+\Delta t} \tag{5.2}$$

This re-balancing step is repeated until the desired future date. Third, on that final date, the portfolio is unwound and the hedging error is computed using formula (5.2).

Our different criterion - payoff replication - presents two attractive features:

- First, instead of relying on the discrete formula for the hedging error as given by equation (5.2), we rely on the continuous process for the replicating portfolio given in equation (3.9) to calculate the replication error (given by the difference between the replicating portfolio value at maturity and the final payoff). This means that cross-series of option data are not required to compute payoff replication errors since, in addition to the stock price process and the payoff, the replicating portfolio can also be simulated under the different models. We are therefore able to analyze hedging performance both theoretically and empirically. This will enable to study the over-hedging behaviour based on HN simulated data as detailed in section 5.4. We stress that the procedure used in the literature could also be simulated theoretically, but that the authors don't specify a process for the Delta-neutral portfolio and only resort to empirical analyses.
- The out-of-sample hedging errors are calculated on small time periods in the literature: one day and five days in Bakshi et al. (1997) and Yung and Zhang (2003), seven days in Dumas et al. (1998) for instance. Conversely, by construction, our hedging criterion accounts for the whole life of the option. Since parameters will (at first) be kept constant during the life of the option, our measure represents a more challenging test to compare the hedging performance of our models.

In short, with payoff replication, we will be able to better understand the over-hedging effect from a theoretical point of view and, empirically, we will manage to appraise whether or not our measure leads to different results than those found in the literature: a model that hedges accurately on a few days horizon may not automatically be accurate until maturity.

5.3 Out-of-Sample Payoff Replication Methodology

In this section, we explain how we proceed to analyze the payoff replication performance based on the data described in section 5.1. This analysis is done out-of-sample since this is more relevant for practitioners than in-sample. The four upcoming sections explain the market calibration procedure used to estimate the parameters, the initial value, lower and upper bounds of the parameters, our sliding window technique and finally the error measure used to classify the models.

5.3.1 Market Calibration Procedure

Chapters 3 and 4 explained that each model will be calibrated to market data to estimate the parameters based on a pricing $MSE loss function. The calibrations are done on Matlab with the *lsqnonlin* (least-squares non-linear) function, which is a local optimizer. We want to stress an important point regarding our optimizations. Because we are dealing with non-linear models, the calibration may "present the problem that the objective function is not necessarily convex and may exhibit several local minima" (Crisostomo, 2014). In particular, the optimal solution obtained might be very dependent on the initial guess of the parameter set. This is especially relevant for HN and Heston where five parameters are estimated. To counter this problem, one solution could be to use a global optimizer such as *asamin* (a simulated annealing algorithm). However, the resulting computational times are substantially increased and, also, Crisostomo (2014) found that, for the Heston model, local optimization with *lsqnonlin* provided better calibration than *asamin*. Still, the main drawback of local optimizers is that we can not be sure whether the solution found is the best one, and we can not measure how far it is from the best one. To ensure that we are not too far from the global minimum, it is therefore necessary to carefully choose the initial values as well as lower and upper bounds of the parameters.

5.3.2 Initial Values & Bounds of Parameters

To ensure that the optimal parameters' values make sense from an economic point of view and that we are not too far from the global minimum of the $MSE loss function, we use realistic initial values as well as lower and upper bounds for the optimization.

Black-Scholes Following Heston and Nandi (2000), the initial value of the volatility parameter σ is set equal to the one-year historical volatility of log-returns of the time series. The lower and upper bounds are fixed to 0% and 100%, a volatility higher than 100% being unrealistic, even in crisis periods.

Practitioner Black-Scholes The initial value of θ_0 is fixed to 1, θ_1 to -1 (a negative value to account for the skew) and θ_2 to 1 (a positive value to account for the smile). The PBS model optimization is not very complex, so that we can choose large bounds for the parameters: the upper and lower bounds are [-50,50] for the three parameters, which gives a wide range of possible values.

Heston-Nandi For this model, finding initial values leading to a low pricing $MSE was not as easy as for the other models. After several trials, we found that the following initial conditions lead to good results: h_0 is equal to the one year historical variance of log returns (transformed into a daily variance), $\beta_0 = 0.0001$, $\beta_1 = 0.0001$, $\omega = 150$ and β_2 is fixed through the persistence formula: $\beta_1 + \beta_2\omega^2 = 0.85$ ($\Leftrightarrow \beta_2 = 3.78 * 10^{-5}$). The upper and lower bounds used are: $h_0 \in$ [0,1/253] β_0, $\beta_1 \in$ [0,1], $\omega \in$ [0,1000] and $\beta_1 + \beta_2\omega^2 \in$ [0.4,1]. The lower bound for β_0 and β_1 and upper bound for $\beta_1 + \beta_2\omega^2$ ensure the stationarity of the model.

Heston The parameters' initial values are: V_0 and θ are equal to the one year variance of log-returns, $\alpha = 0.2$, $\rho = -0.5$ and κ is set as a byproduct of the Feller condition: $2\kappa\theta - \alpha^2 = 0.1$ ($\Leftrightarrow \kappa = 1.75$ if $\theta = 0.2^2$ for example). The upper and lower bounds used are: V_0, $\theta \in$ [0,1], $\alpha \in$ [0,5], $\rho \in$ [-1,1], $2\kappa\theta - \alpha^2 \in$ [0,20].

In order to consider time-consistency in the parameters, we also compare the $MSE under these initial values with the $MSE obtained by using the estimated parameters of the last window as initial values (see section 5.3.3 for a description of the sliding window technique), and we choose the optimal parameters accordingly.

5.3.3 Sliding Window Technique

To conduct our payoff replication performance analysis on real data, which is done out-of-sample, we use a *sliding window technique*. The size of the window is two trading days. This means that, first, we estimate the parameters over the first two days, and then that we shift the window by two days and perform the payoff replication on the options over these next two days.[6] The window is then shifted again by two days and the parameters are re-estimated. This means that half of the windows are used for parameters' estimation purposes and the other half to calculate the out-of-sample replication errors, for a total of 30 windows for each time series.

Practically, we use the parameters estimated in the current time window to calculate the hedging errors for the next time window by using the time series of S&P500 or Apple prices. Each day, the daily value of the Delta and thus of the replicating portfolio are calculated, until expiration of the option where the replication error is computed as

$$\varepsilon_T = \Pi_T - \max(S_T - K, 0) \tag{5.3}$$

where Π_T is the value of the replicating portfolio at maturity. The replicating portfolio continuous process in equation (3.9) has to be discretized and, as mentioned by Yung and Zhang (2003), when the underlying distributes dividends, the formula must be adjusted accordingly, becoming:

$$\Pi_{t+\Delta t} = \Pi_t + \Delta_t(S_{t+\Delta t} - S_t) + r\Delta t(\Pi_t - \Delta_t S_t) + \Delta_t D_{\Delta t} \tag{5.4}$$

where $D_{\Delta t}$ is the daily dividend from the underlying asset during the period Δt.[7] However, as noted in section 5.1, the S&P500 is already adjusted to dividends, so that the extra term $\Delta_t D_{\Delta t}$ will not be added. It is however added for Apple because we use the closing prices and not the adjusted ones that account for the extra return from dividends. Importantly, following Petitjean and Moyaert (2011), we take the initial value of the replicating portfolio to be the market price of the option, i.e. $\Pi_0 = C_0^{mkt}$, "so that the results are not affected by the difference in pricing between the models." We take 253 trading days in a year since 2016 is a leap year.

For the S&P500 time series, on average, 95 option contracts were used in each estimation window, and 73 in each hedging window. For Apple, it amounts to 89 and 81 option

[6] Remember from section 5.1 that, for the windows used for hedging purposes, some options had maturity dates extending beyond the time of the analysis and therefore had to be eliminated.

[7] The Delta in equation (5.4) is calculated after subtraction of the discounted present value of dividends from the underlying asset price. Conversely, S_t in equation (5.4) is not adjusted to future dividends. The reason is that, in practice, the portfolio is re-balanced by buying or selling the underlying asset at the market price itself.

contracts respectively. A detailed breakdown of the number of contracts used in the estimation and hedging windows can be found on tables C.2 and C.3 in appendix.

Using such a procedure represents a reasonable trade-off between the sample size in each time window that must be large enough for estimation purposes, and the time dependence of the parameters that calls for frequent recalibration. Moreover, using only daily calibration, as explained by Bates (2003), can hide fundamental flaws in the underlying model, which is an issue that we wanted to avoid.

5.3.4 Error Measures & Classification of Models

The error measures used to appraise the distance between the replicating portfolio value and the payoff at maturity is the Dollar Mean Absolute Error (\$MAE) for ease of interpretation. In order to capture any potential bias (under-hedging or over-hedging), we also use the normalized Dollar Mean Error (normalized \$ME). The normalized \$ME is calculated as the \$ME divided by the \$MAE. The sign indicates whether the model under- or over-hedges, and the closer it is to 1 in absolute value, the stronger is the bias. The \$MAE and normalized \$ME are reported on our moneyness-maturity categories for each model and each time series, and are calculated as, with n the number of options:

$$\begin{aligned} \$\text{MAE} &= \frac{1}{n}\sum_{i=1}^{n} \left|\Pi_{T,i} - \max(S_{T,i} - K_i, 0)\right| \\ \$\text{ME} &= \frac{1}{n}\sum_{i=1}^{n} \Pi_{T,i} - \max(S_{T,i} - K_i, 0) \end{aligned} \tag{5.5}$$

We wish to stress that we will take the best model as the one having the lowest \$MAE. However, one may wonder why we do not use a maximum \$ME criterion given that over-hedging the payoff at maturity is attractive since you are left with a positive cash bonus, while conversely under-hedging the payoff is not desirable. However, in practice, the goal of hedging is not to maximize the value of the replicating portfolio but to have a hedge as accurate as possible. Moreover, we use a simple Delta-neutral portfolio whose Delta, and thus Gamma, is positive. However, in practice, practitioners have a whole portfolio of options and may in fact have a short position in some options and a long position in others, and therefore different relations to Gamma when hedging them.[8] As pointed out by Joshi (2008), this has an important effect: being short Gamma means that the procedure of hedging will cost money over the life of the option, and the contrary if we are long Gamma. Our empirical results in the next chapter will confirm that we indeed observe an over-hedging effect in average given that our replicating portfolio has a positive Gamma. In brief, over-hedging is not always a positive feature since it will transform into under-hedging if one holds a Delta-neutral portfolio having a negative Gamma. This is why we consider the best model to be the one as close as possible to the value of the payoff at maturity.

5.4 Payoff replication on Simulated Data Methodology

To provide explanations to the over-hedging effect observed on our empirical data, we rely on a simulation analysis. This means that we simulate ourselves the stock price dynamics under carefully chosen parameters' values to calculate the payoff replication errors. The advantage is that we place ourselves in a "perfect" experiment where we can freely control the value of the parameters whose effect we want to understand. Conversely, with empirical data, you don't have access to such a flexibility and several effects interact with each other (e.g. the estimation method, the underlying asset behaviour, imperfect models), which make it harder to understand the impact of different variables on payoff replication behaviour, *all else being equal*.

The simulations are done on HN because its discrete-time form lends itself well to test the impact of the re-balancing frequency on under/over-hedging. Practically, we re-balance the replicating portfolio under the HN stock price dynamics with the HN Delta until expiration of the option. Since we are not in a pricing perspective, we use the processes under the physical measure and not the risk-neutral one. This requires to be careful to the relations between the physical and risk-neutral parameters. Indeed, the call option premium under HN at $t = 0$ and the HN Delta are calculated under the risk-neutral

[8] Moreover, hedging put options requires the opposite position in the underlying asset than when hedging calls, which makes that the sign of the Gamma for the Delta-neutral portfolio when hedging a put will be opposite to the one for a call.

measure, while the HN stock price dynamics is simulated until $t = T$ under the physical measure. The risk-neutral and physical processes were presented in section 4.2.2.

We distinct three levels of volatility ($\sigma = 15\%, 30\%, 45\%$), four re-balancing frequencies ($k = 0.5, 1, 2, 5$, i.e. twice a day, daily, every two days and every five days) and finally four levels of expected return ($\mu = -10\%, -5\%, 5\%, 10\%$). To fix the volatility in the context of the HN model, what we have chosen to do is that when we have $\sigma = 15\%$ for example, we fix the initial annualized variance ($252 * k * h_0$) and the long-term variance ($252 * k * \frac{\beta_0+\beta_2}{1-\beta_1-\beta_2\gamma^2}$) to $(15\%)^2 = 2.25\%$. For the other parameters, we choose plausible values as in figure 4.3: we set $r = 3\%$, $\beta_0 = 8 * 10^{-6}$, $\beta_1 = 0.45$, $\gamma = 700$, β_2 is fixed to fulfill the long-term variance condition, and λ is chosen such that $\mu = r + 252 * k * \lambda * h_0$. Finally, we fix $T = 0.5$ and $S_0 = K = 50$, i.e. at-the-money options.

The numerical procedure that we use is Monte Carlo simulation with 1000 random paths.[9] Since we are assessing under or over-hedging, we only report the normalized \$ME for the different levels of volatility, expected return and re-balancing frequency.

[9]This is lower than the 5000 random paths recommended by Christoffersen (2012) because HN is quite computationally intensive. We also performed some calculations with 5000 random paths to assess whether the differences in results were substantial or not, and found that the results were very close so that the conclusions were the same.

Chapter 6

Payoff Replication: Results & Further Analysis

Based on the methodology detailed in chapter 5, we now report and analyze our results to answer to our research questions. First, we provide statistical descriptions of our two time series and a short analysis of the calibrated parameters. Then, we describe the payoff replication findings on our two time series. Additionally, we provide two further analyses: we compare the results with the out-of-sample pricing performance and we look at the effect of a frequent re-calibration of parameters. Finally, we provide explanations for the observed over-hedging effect based on HN simulations.

6.1 Time Series Statistical Description

This section describes the time series of log-returns of the S&P500 and Apple over the time period considered. We particularly want to assess whether the data detract from the BS description of normally distributed log-returns. This analysis is of interest because, if GBM proves to represent an accurate description of our data, then we could expect BS to outperform the PBS, HN and Heston models.

S&P500 Figure 6.1 shows that the log-returns fit quite well to the Normal density. The normality test of Jarque and Bera (1987) confirms that normality can not be rejected with a p-value of 79.03%. Then, figure B.3 in appendix shows that log-returns are not auto-correlated. This is in line with the models used in this thesis that all assume log-returns to present no auto-correlation. Finally, figure B.5 in appendix shows strong evidence of auto-correlated squared log-returns, i.e. of volatility clustering, which is explicitly modeled by the HN and Heston models.

FIGURE 6.1: Histogram of log-returns of the S&P500 index from January 4 to May 31, 2016 compared to the normal density.

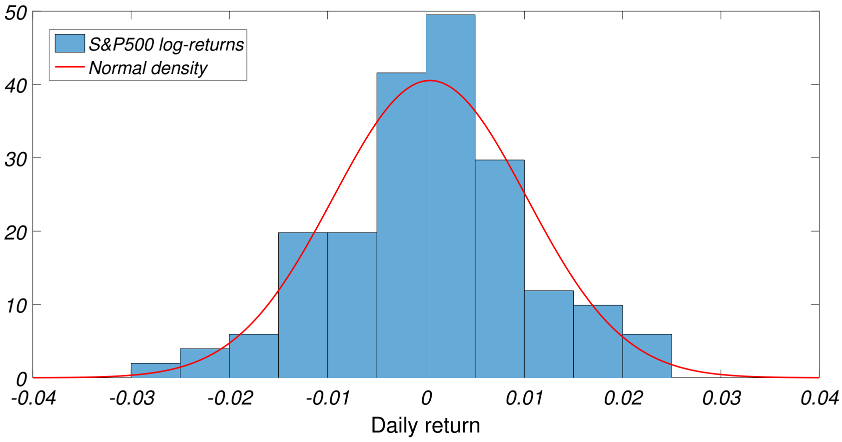

Notes. The time span goes from January 4, 2016, which is the first date of our sample of options, to the largest maturity considered for payoff replication, i.e. May 31, 2016.

FIGURE 6.2: Histogram of log-returns of Apple stock from January 4 to June 17, 2016 compared to the normal density.

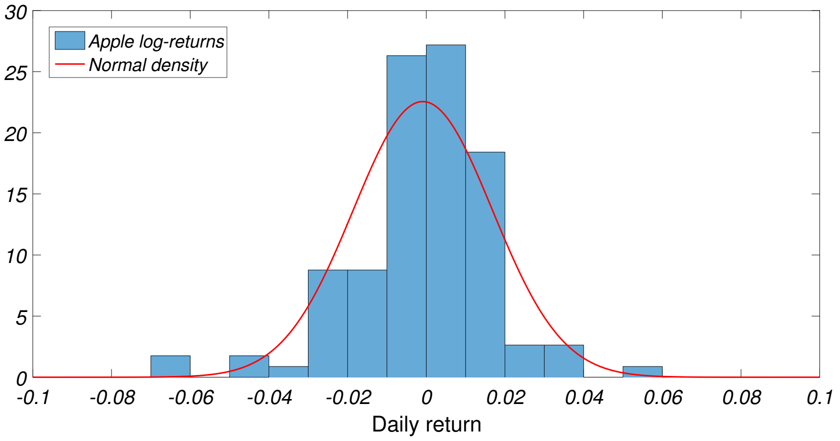

Notes. The time span goes from January 4, 2016, which is the first date of our sample of options, to the largest maturity considered for payoff replication, i.e. June 17, 2016.

Apple Figure 6.2 shows that Apple diverges quite strongly from the Gaussian. This is confirmed by a p-value of $3.51 * 10^{-11}$ from the Jarque-Bera normality test.[1] Then, figure B.4 in appendix shows that log-returns do not depict auto-correlation. Finally, figure B.6 in appendix shows no evidence of volatility clustering.

All in all, whether due to non-normality (Apple) or volatility clustering (S&P500), we should except BS to be outperformed by the other three models since the geometric Brownian motion is a poor description of the two assets. In fact, our results will show

[1]The skewness equals -0.7269 and the kurtosis 5.8323, compared with -0.1143 and 3.2442 for the S&P500.

that BS performs worst for *pricing* but not for *hedging* purposes. This is due to the fact that we use a pricing \$MSE estimation loss function while optimizing the fit to option prices does not automatically imply accurate Delta estimates. As a result, BS will very often beat HN and Heston (but not PBS) because these two models tend to over-fit the option price curvature and are therefore not accurate for hedging.

6.2 Parameters Analysis & Discussion

Plots of the estimated parameters across the 15 estimation windows are displayed in appendix D. Tables D.1 and D.2 in appendix report summary statistics of the parameters. By looking at the ratio between the standard deviation and the mean of parameters, we see that, for the S&P500, PBS has the most varying parameters with all ratios being above one, while it is significantly more stable for Apple. However, figures D.2 and D.6 in appendix show that θ_1 decreases when θ_0 and θ_2 increase (and vice versa), so that the resulting implied volatilities are quite stable through time. BS implied volatility is also stable for both time series. Conversely, the parameters of HN and Heston often present substantial time-variations, which indicates the relevance of having a frequent re-estimation of the parameters for these two models.

We wish to point out that large time-variations of parameters, such as for HN and Heston, indicate that the models do not always accurately represent the data dynamics. Indeed, a good model can incorporate information on the underlying asset distribution with sufficiently stable parameters through time. Conversely, a model with unstable parameters may still be able to accurately fit the option price curvature with a sufficient number of *degrees of freedom* (parameters) but at the cost of a frequent re-calibration of the parameters. In that case, it can either mean that the parameters are hard to estimate or that they don't represent their intended financial/economic meaning.[2]

[2]Examples of seemingly nonsensical values are: ρ is very close to -1 under the S&P500, the volatility of volatility α takes very high values for both time series, even often higher than 100% for Apple, and V_0 and h_0 are sometimes absurdly close to 0 for Apple.

6.3 Out-of-Sample Payoff Replication Results

In this section, we report and interpret the payoff replication performance of our four models based on the data described in section 5.1 and the methodology of section 5.3.

6.3.1 Analysis of Out-of-Sample Payoff Replication

This section answers to our main research question: *Among the Black-Scholes, Practitioner Black-Scholes, Heston-Nandi and Heston models, which one performs best in terms of out-of-sample hedging performance, as measured by payoff replication, and how does performance differ according to the level of moneyness, maturity and volatility of the underlying?*. We also comment on the observed under/over-hedging bias.

The out-of-sample payoff replication \$MAE are reported in table E.1 in appendix, while the normalized payoff replication \$ME are reported in table E.2. We first distinguish the results for the two time series and then end up by looking at the ranking of models in function of the end-moneyness rather than the moneyness at inception.

S&P500 Index

Over-hedging bias. First, the normalized \$ME criterion shows that all models over-hedge on average. This bias increases with maturity, except for DITM options. It also generally decreases when moneyness rises, meaning that OTM options lead to more over-hedging than ITM ones. Our results diverge from Bakshi et al. (1997) and Yung and Zhang (2003) who found, by using S&P500 option quotes, that the mean hedging errors for the BS and SV models (for the first authors and on call options) and PBS and EGARCH models (for the second author and on put options) were systematically negative. Section 6.4 will explain that it is likely due to the fact that we hedge until expiration while the authors only considered one and five days horizons. Section 6.4 will also provide theoretical explanations on simulated data to this over-hedging behaviour, which is linked to the positive Gamma as mentioned in section 5.3.4.

Effect of maturity. Table E.1 as well as figure 6.3 show that the \$MAE increases with maturity for OTM and ITM options, while it decreases for DITM ones (except for PBS). For DOTM options, the errors slightly decrease with maturity for BS and PBS, while

Figure 6.3: Out-of-sample payoff replication $MAE for mid-term and long-term options for the S&P500 index.

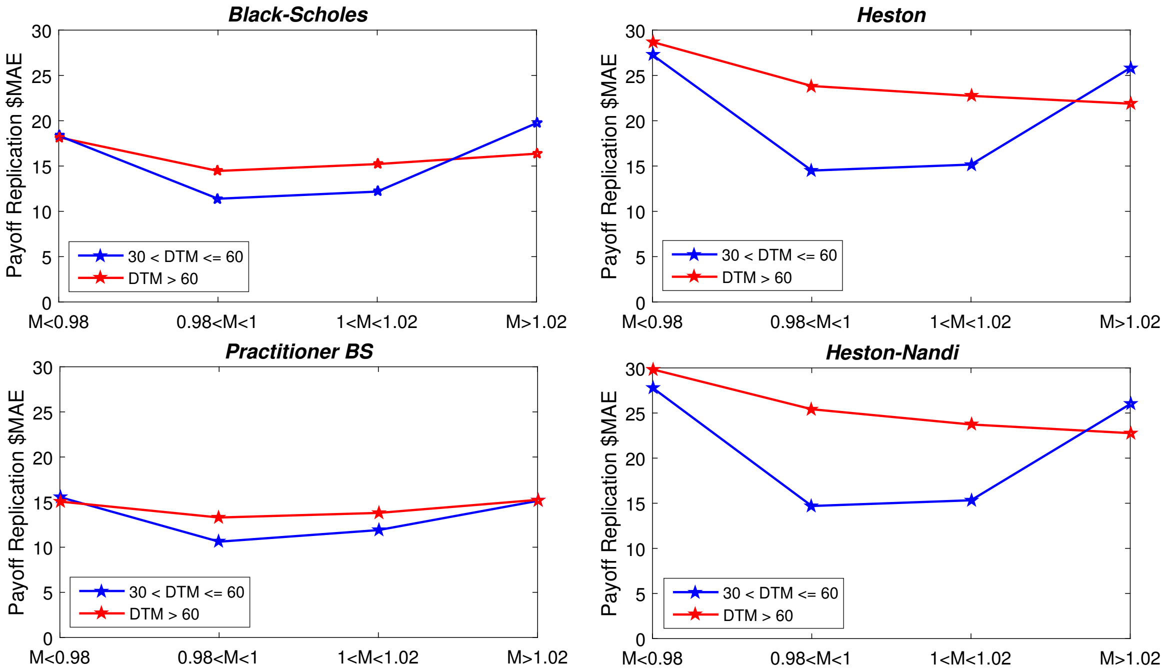

Notes. The S&P500 options cover the period January 4, 2016 to March 30, 2016. M on the x-axis refers to moneyness, defined as $\frac{S_t - D_{pv}}{K}$, and DTM to Days To Maturity. We display the $MAE for each moneyness-maturity category for mid-term and long-term contracts. The errors are calculated out-of-sample on each of our hedging windows, which are composed of two days, by using the parameters estimated on the previous estimation window, also composed of two days. 15 hedging windows are used in total over the sample period.

they increase for HN and Heston. The fact that the $MAE sometimes decreases with maturity is not intuitive since one could reasonably except payoff replication to be less accurate when you have to hedge during a longer period of time, especially given that the parameters are kept constant until expiration of the option.

Effect of moneyness. The effect of moneyness (at inception) is consistent across all models and depends on maturity. For short-term options, table E.1 shows that moneyness has a very minor effect (since there is less time-uncertainty): the $MAE slightly decreases when going from OTM to ITM options, expect for Heston. For mid-term options, the errors are lower for OTM and ITM options than DOTM and DITM ones, which results in the U-shaped blue curve in figure 6.3. Finally, for long-term contracts, the errors decrease with moneyness for HN and Heston, while we once again have a U-shaped curve for BS and PBS. Overall, the results outline that near-the-money options are generally more accurately hedged than far-from-the-money ones.

Figure 6.4: Graphical ranking of models across moneyness and maturity based on the out-of-sample payoff replication \$MAE for the S&P500 index.

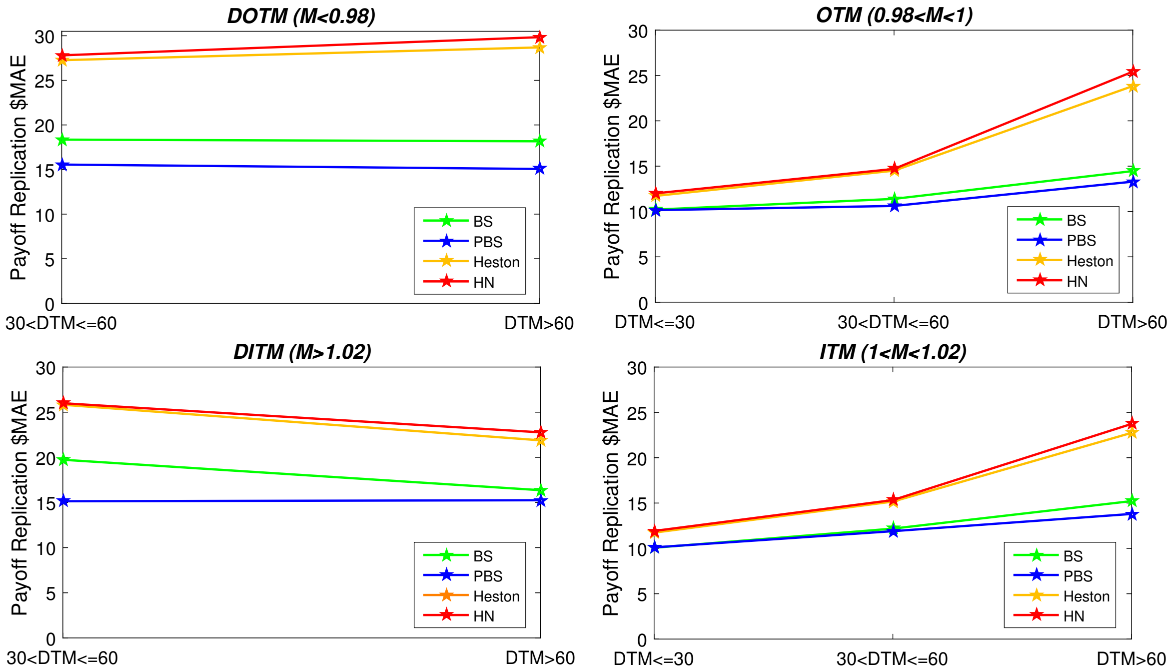

Notes. The S&P500 options cover the period January 4, 2016 to March 30, 2016. M refers to moneyness, defined as $\frac{S_t - D_{pv}}{K}$, and DTM to Days To Maturity. We display the \$MAE for our different levels of moneyness and maturity to assess the ranking between the four models. The errors are calculated out-of-sample on each of our hedging windows, which are composed of two days, by using the parameters estimated on the previous estimation window, also composed of two days. 15 hedging windows are used in total over the sample period.

Ranking of models. Figure 6.4 establishes a very clear ranking among the four models. For all categories, Heston and HN are largely outperformed by BS and PBS. Heston also always slightly outperforms HN. Considering that HN, as explained in section 4.3.4, is a discrete-time version of Heston, we see that discretizing time does not enable to achieve a better hedging performance. This result disagrees with Petitjean and Moyaert (2011) who, as mentioned in section 4.2.5, had found lower hedging errors for HN than Heston during the subprime crisis. We will see afterwards that the results on the more volatile Apple stock will not diverge as much from the authors' findings. Regarding BS and PBS, PBS always outperforms BS, therefore systematically being the best of our four models for payoff replication. This disagrees with Dumas et al. (1998) who found that PBS performs less well for hedging than BS. This may be due to the different functional form that we use (see equation (4.2)).

Although we just concluded that Heston and HN are systematically the worst performers in all moneyness-maturity categories, figure 6.4 also shows that the ranking between the four models is least striking for short-term contracts. This indicates that Heston and HN could perform better than BS and PBS for very short-term options. To verify this, we report the payoff replication $MAE for options with 20, 15 and 10 or less days to maturity in table E.3 in appendix. The results show that our previous findings are robust to a short maturity: the ranking between the four models stays the same. However, the errors are quite close to one another because, when maturity is short, the Delta quickly converges to values very close to 0 or 1 (depending on whether the option is OTM or ITM) and so the model specification has less impact for those options.

Over-fitting issue. All in all, the results agree with the findings of Baskhi et al. (1997), Dumas et al. (1998) and Yung and Zhang (2003) who found that more complex model specifications than the simple BS or PBS models do not improve or even decrease the hedging performance. More advanced volatility modelling as in the HN and Heston models may thus produce superior pricing performance, as per the literature findings and as section 6.3.2 will show, but do it at the cost of more parameters to estimate. This creates issues of over-fitting and therefore less stable option valuation errors through time, which, as pointed out by Dumas et al. (1998), lead to higher hedging errors. In other words: "[...] more parameters in a model may result in more freedom in fitting option price curvature, but may generate less accurate delta estimates" (Yung and Zhang, 2003). Given that we keep the parameters constant until maturity, the parameters instability (from the over-fitting issue) outlined in section 6.2 for HN and Heston plays a great role. Interestingly, this also points out that HN and Heston may be able to perform better at the cost of re-estimating the parameters sufficiently often throughout the life of the option. Section 6.3.3 will provide such an analysis and will show that the ranking is not altered and even that the errors will increase: a frequent re-calibration does not help in generating more accurate Delta estimates.

Apple Stock

We now report the results for the more volatile Apple stock and compare them with those uncovered for the S&P500. We want to appraise whether a larger volatility of the underlying asset changes some of our findings, keeping in mind that the two time series do not only distinguish by their volatility as section 6.1 showed.

FIGURE 6.5: Out-of-sample payoff replication $MAE for Apple stock.

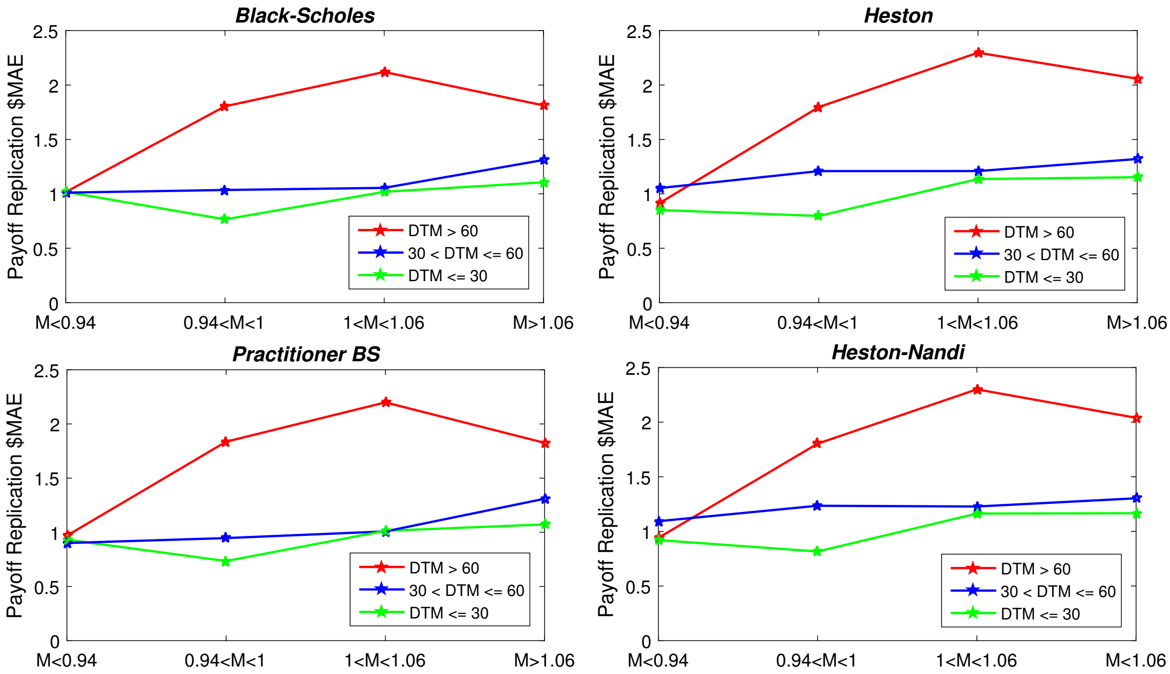

Notes. The Apple stock options cover the period January 4, 2016 to March 30, 2016. M on the x-axis refers to moneyness, defined as $\frac{S_t - D_{pv}}{K}$, and DTM to Days To Maturity. We display the $MAE for each moneyness-maturity category. The errors are calculated out-of-sample on each of our hedging windows, which are composed of two days, by using the parameters estimated on the previous estimation window, also composed of two days. 15 hedging windows are used in total over the sample period.

Over-hedging bias. Table E.2 in appendix shows that all the models systematically over-hedge, as we had found for the S&P500. The normalized $ME nearly always increases with maturity. Compared to the S&P500, PBS systematically over-hedges more for Apple, and also quite often the BS model, while the normalized $ME tends to be lower under HN and Heston. Across all models, there is systematically more over-hedging for Apple than the S&P500 for short-term options.

Effect of maturity. From figure 6.5, we see that the errors are overall increasing with maturity, except for DOTM options where long-term options lead to very close or lower errors than short-term and mid-term ones.

Effect of moneyness. Third, figure 6.5 and table E.1 show that the effect of moneyness for different maturities is once again very consistent across all models. The most striking feature that we want to outline appears for long-term contracts: we find a reverse U-shaped curve, meaning that the errors first largely increase with moneyness and then also quite strongly decrease from ITM to DITM options. This shape is at odds with what we had found for the S&P500. This indicates that the volatility of the underlying (and its whole distribution more generally) strongly impacts how moneyness affects the replication errors, and can even give opposite effects.

Ranking of models. It is interesting to look at whether HN and Heston are still outperformed under the Apple stock. The ranking is not as straightforward as for the S&P500. Table E.4 in appendix reports the ranking of the models in each moneyness-maturity category. Figure 6.6 displays the ranking graphically. The first thing that we observe is that the models display closer results than under the S&P500. Then, for ITM and DITM options, BS and PBS still outperform HN and Heston. PBS slightly surpasses BS for short and mid-term options, and the contrary for long-term ones. This contrasts with the results on the S&P500 where BS was always beaten by PBS. Interestingly, divergent results compared to the S&P500 are found for OTM and DOTM options: HN and Heston are now the two best models for long-term contracts and DOTM short-term contracts, with Heston being first, while otherwise they are outperformed by BS and PBS, with PBS being first. These results only partially confirm the finding of Petitjean and Moyaert (2011) that Heston outperformed BS for hedging during the very volatile subprime crisis. HN is still not a good alternative to Heston for hedging purposes, though they generate very close results. Overall, considering that larger moneyness levels are available for Apple than for the S&P500, our results demonstrate that under a large enough volatility, more complex models such as Heston and HN are able to compete with and even outperform simpler models for options that are sufficiently OTM. This is in line with Bakshi et al. (1997) who had found, as mentioned in section 4.3.5, that the only real improvement brought by SV models over BS was for OTM calls. In fact, we will later show, both for the S&P500 and Apple, that HN and Heston perform well for options that are sufficiently OTM *at maturity*, i.e. when the Delta is close to 0 during a large proportion of the hedging period.

Effect of volatility. Finally, we wish to assess whether a higher volatility of the underlying asset implies higher payoff replication errors. To compare the errors under the S&P500 and Apple, table E.5 reports the normalized replication $MAE, i.e. the mean of

FIGURE 6.6: Graphical ranking of models across moneyness and maturity based on the out-of-sample payoff replication $MAE for Apple stock.

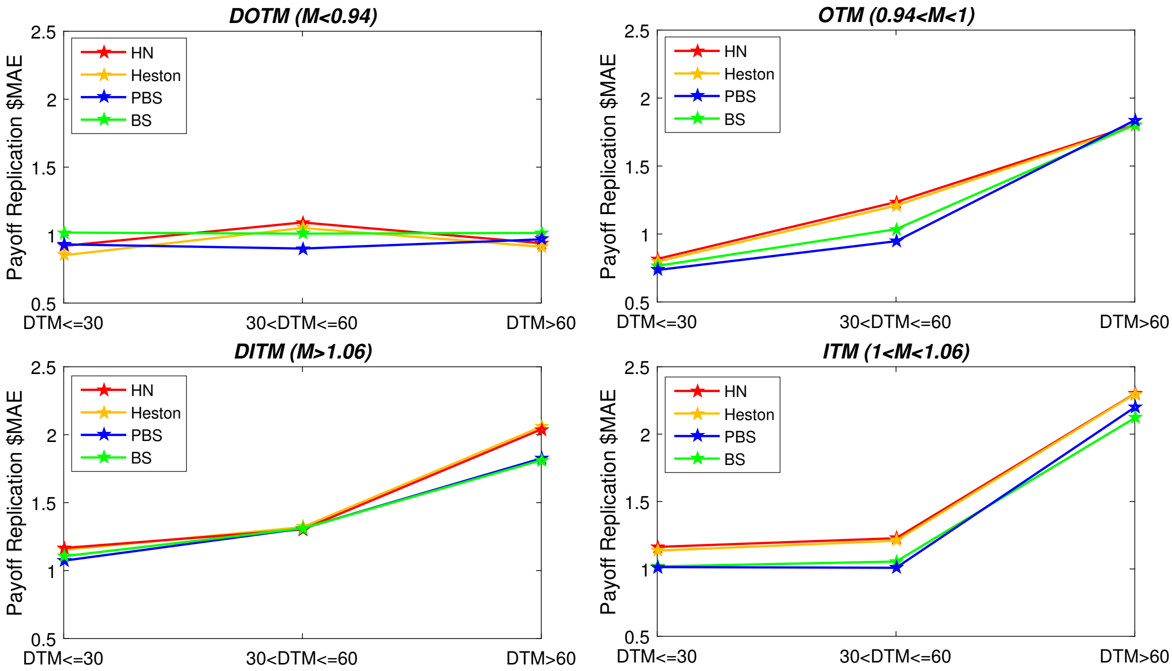

Notes. The Apple stock options cover the period January 4, 2016 to March 30, 2016. M refers to moneyness, defined as $\frac{S_t - D_{pv}}{K}$, and DTM to Days To Maturity. We display the $MAE for our different levels of moneyness and maturity to assess the ranking between the four models. The errors are calculated out-of-sample on each of our hedging windows, which are composed of two days, by using the parameters estimated on the previous estimation window, also composed of two days. 15 hedging windows are used over the sample period.

absolute errors divided by the initial call option market price C_0^{mkt}.[3] Table E.6 reports the errors averaged by maturity and moneyness. In average across all categories[4], the errors are quite largely lower for the S&P500 when hedging with BS and PBS, while they are larger under HN and Heston, which shows that these two models are better suited for more volatile underlying assets. When dividing by maturity, we in fact notice that HN and Heston lead to lower errors for Apple especially for long-term maturities. Across moneyness, lower errors for Apple do interestingly emerge under BS and PBS for DITM options, while for Heston and HN it appears for both DOTM and DITM options. These results overall indicate that more stable underlying asset returns do not always induce more accurate hedging, which is not especially intuitive.

[3]The absolute hedging errors are indeed substantially higher for the S&P500 than for Apple because the average call option price over the total sample is also very much greater: $62.68 for the S&P500 against $7.12 for Apple.

[4]Not accounting for short-term DOTM and DITM contracts that are not available for the S&P500.

End-Moneyness Ranking

The ranking between the four models has been established in function of the moneyness at inception. However, being ITM or OTM at $t = 0$ is of course no guarantee that it will also be the case at $t = T$. In particular, from figure 5.1, both the S&P500 and Apple have had periods of high surge and decline during the time span considered that impacted the end-moneyness. Looking at the ranking in function of the end-moneyness, i.e. $\frac{S_T}{K}$, may change some of the previous conclusions. Indeed, hedging options that finished DOTM or DITM means that it is likely that the Delta was close to 0 or 1 during a substantial part of the hedging period. In turn, the model specifications have less impact for those options than those finishing near-the-money.

The number of options used in each maturity and end-moneyness category for the two time series is displayed on table E.7 in appendix.[5] The payoff replication \$MAE for our two time series in function of the end-moneyness are reported on table E.8, along with graphical displays of the rankings on figures 6.7 and 6.8. The most striking observation is that HN and Heston compare favourably to BS and PBS for options that finish OTM. Specifically, for the S&P500, Heston and HN finish first and second for options that have a moneyness below 0.98 at maturity, and very drastically so for a long enough maturity. This diverges from the ranking in function of the moneyness at inception. For Apple, as previously found, Heston and HN outperform BS and PBS for long-term options that finish with a null payoff.

To sum up, more realistic volatility specifications achieve a better hedging when the Delta converges close to 0 during a large enough proportion of the hedging period.

6.3.2 Comparison with Out-of-Sample Pricing Performance

This section answers to our first research sub-question: *Is an accurate fit of the market option prices by a certain model a good indicator of its hedging performance?*

The previous section showed that more complex option pricing models do not compare well to simpler models such as BS and PBS for hedging purposes. We argued that, as exposed by Dumas et al. (1998), this is due to over-fitting which, although can give lower

[5] The definition of moneyness categories is modified. For the S&P500, it is: M<0.98, 0.98<M<1.03, 1.03<M<1.08 and M>1.08. For Apple, it is: M<0.85, 0.85<M<1, 1<M<1.15 and M>1.15.

FIGURE 6.7: Graphical ranking of models in function of the end-moneyness based on the out-of-sample payoff replication \$MAE for the S&P500 index.

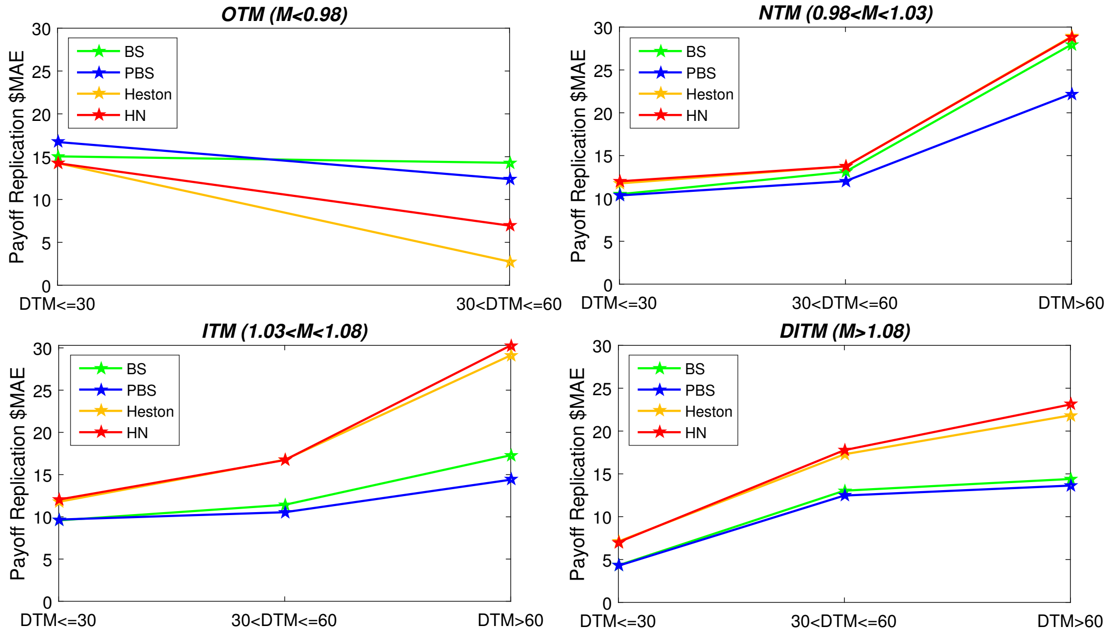

Notes. The S&P500 options cover the period January 4, 2016 to March 30, 2016. M refers to end-moneyness, defined as $\frac{S_T}{K}$, and DTM to Days To Maturity. The errors are calculated out-of-sample on each of our hedging windows, which are composed of two days, by using the parameters estimated on the previous estimation window, also composed of two days. 15 hedging windows are used in total over the sample period.

FIGURE 6.8: Graphical ranking of models in function of the end-moneyness based on the out-of-sample payoff replication \$MAE for Apple stock.

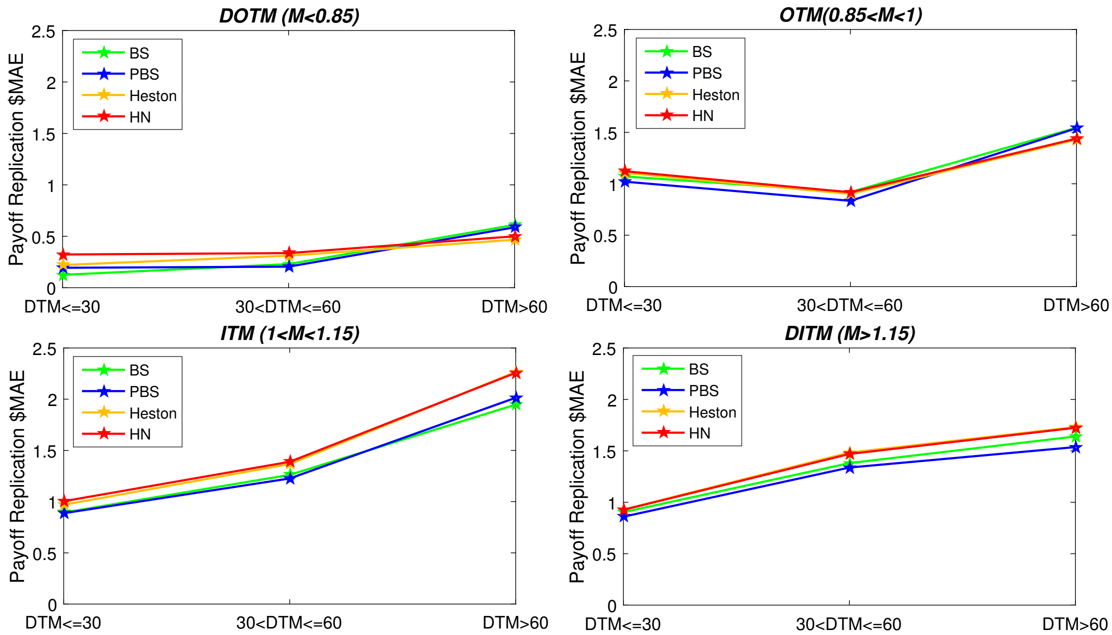

Notes. The Apple stock options cover the period January 4, 2016 to March 30, 2016. M refers to end-moneyness, defined as $\frac{S_T}{K}$, and DTM to Days To Maturity. The errors are calculated out-of-sample on each of our hedging windows, which are composed of two days, by using the parameters estimated on the previous estimation window, also composed of two days. 15 hedging windows are used in total over the sample period.

pricing errors, is a problem for hedging. We now want to explicitly verify whether a better fit of the option price curvature implies less accurate Delta estimates or not.

This section first briefly comments the out-of-sample pricing errors for our four models and two time series.[6] Table F.2 in appendix reports the out-of-sample pricing root-mean-square errors ($RMSE).[7] The number of contracts that were used in each moneyness-maturity category for out-of-sample pricing can be found on table F.1 in appendix. Tables F.3 and F.4 report the ranking of the four models in each moneyness-maturity category for both time series. The comments we can make are, in short:

S&P500 HN and Heston are overall the best performers: they have the two lowest $RMSE in 7 out of the 10 moneyness-maturity categories. HN fares better than Heston by beating it in 8 categories out of 10. BS is a very poor performer: it finishes in the last position in every category except for ITM short-term contracts. PBS does not behave so poorly, especially for DITM and DOTM options where its RMSE closely follows or does better than HN and Heston. The most striking improvement in the $RMSE brought by HN and Heston compared to PBS is for short-term options: PBS performs 43.84% and 32.65% worse than the best model for OTM and ITM options.

Apple This time, Heston performs better than HN for short-term and mid-term contracts, while HN is superior for long-term contracts. We once again observe that BS performs very badly by finishing in the last place in 9 categories out of 12, often by a large margin. Overall, the best model to use is not always the most complex one. Specifically, PBS performs better than HN and Heston for long-term options and DOTM and DITM mid-term options. That said, the improvement compared to HN or Heston tends to be very low for these options. As for the S&P500, the improvement brought by Heston compared to PBS is most striking for short-term options.

Overall, compared to what section 4.1.2 explained that Christoffersen and Jacobs (2003) had concluded from their research, PBS does not compare so well to Heston for pricing

[6]The errors are calculated on the same windows than for the out-of-sample payoff replication errors. We have re-included the options that had been eliminated because they had maturity dates extending beyond the date of analysis (as explained in section 5.1).

[7]The $RMSE criterion is used in order to have consistent estimation and evaluation loss functions, since the parameters were estimated by minimizing the pricing $MSE. The requirement of having consistent estimation and model evaluation loss functions to achieve an optimal out-of-sample loss was pointed out by Christoffersen and Jacobs (2003).

purposes. However, we confirm the results of Heston and Nandi (2000), mentioned in section 4.2.5, who had found that HN outperformed PBS for out-of-sample pricing.

Altogether, these results confirm the conclusions of section 6.3.1 that accurately fitting the market option prices is not a good indicator of hedging performance. Even further, these two measures are at odds with one another. Indeed, HN and Heston fare overall way better for pricing than hedging purposes. The reason is that we use a pricing-based estimation function and that they account for the stylized facts observed in section 6.1.

The relation between hedging and pricing accuracy has been assessed by comparing the models between themselves. Even further, we can wonder if the negative relationship also emerges at the level of a single model: does a low pricing error for a particular option under a certain model implies a large payoff replication error for the same option under the same model? To test this, we run the following regression analysis[8]:

$$\begin{aligned}|\text{Replication Error}|_i = &\beta_0 + \beta_1 * \text{Moneyness}_i + \beta_2 * \text{Moneyness}_i^2 + \beta_3 * \text{Maturity}_i \\ &+ \beta_4 * |\text{Pricing Error}|_i + \varepsilon_i\end{aligned}$$

for $i = 1, ..., n$ where n is the total number of options used for the regression. The maturity is measured in unit of years and moneyness is defined as $\frac{S_t - D_{pv}}{K}$. We are interested in the sign and significance of β_4. The other three explanatory variables have been included to avoid an omitted-variable bias.[9] The results for the S&P500 and Apple are reported on tables G.1 and G.2 in appendix. For the S&P500, β_4 is positive for BS and Heston, and negative for PBS and HN, indicating that a negative relationship at the level of a single level is indeed sometimes observed. That said, none of the β_4 are significantly different from 0 at a confidence level of 5%. The only low p-value observed is for PBS: 8.86%.[10] For Apple, β_4 is positive for all models except for BS where it is slightly negative. This is in line with our finding that more complex models fare better for hedging than under the

[8] The linear regressions are conducted with the statistical analysis software R. The parameters of the regression are estimated by OLS. The homoskedasticity assumption of the error term ε has been tested with the test of Breusch and Pagan (1979), and the assumption of non auto-correlated errors with the Lagrange-multiplier test developed by Breusch (1978) and Godfrey (1978). Both tests rejected the two assumptions for all models with a p-value $< 2.2 * 10^{-16}$. As a result, HAC (Heteroskedasticity and Auto-correlation Consistent) standard deviations of the estimated parameters have been used by resorting to the variance-covariance matrix of Newey and West (1987), with a lag value of 10.

[9] The variable Moneyness^2 has been included given that we previously found U-shaped (for S&P500) and reverse U-shaped (for Apple) relations between the hedging errors and moneyness.

[10] This means that the unilateral test of the form $H0 : \beta_4 = 0$, $H1 : \beta_4 < 0$ for PBS gives a p-value twice as low, i.e. 4.43%, which is below 5%.

S&P500 and are still good for pricing. However, β_4 is only significantly different from 0 for HN with a p-value of 0.32%.

Overall, we see that out-of-sample pricing is not a good predictor of out-of-sample payoff replication at the level of a single model. A negative relationship does not seem to be concluded, except possibly for PBS under the S&P500.[11] This confirms that a better fit of option prices does not especially imply more accurate Delta estimates.

6.3.3 The Effect of a Frequent Re-Calibration of Parameters

This section answers to our second research sub-question: *What is the effect of a frequent re-calibration of parameters on the amplitude of payoff replication errors and on the relative performance of our models compared to keeping parameters constant over the life of option?*

This question is worth investigating since we suggested in section 6.3.1 that Heston and HN may be able to better compete with BS and PBS for payoff replication by applying a frequent re-calibration of parameters. The rationale came to us from Dumas et al. (1998) who explained that more complex models with more parameters create over-fitting and hence larger hedging errors. Since we keep parameters constant and that we hedge until expiration of the option, this effect has an important impact. By re-estimating the parameters sufficiently often, one would be able to cope with the time instability of parameters outlined in section 6.2. As a downside, being constrained to frequently update the parameters to achieve an accurate hedging would also underline that the model has difficulties in representing the data dynamics.

Practically, the Delta is calculated with the parameters of the last estimation window. This means that re-calibration is done every four days (since the estimation windows are separated by two days). The exercise is done on the options within our hedging windows. We eliminate options with maturities extending beyond the final date of our sample of options - March 30, 2016 - because we can't re-calibrate the parameters after that date. As a result, we change the definition of our moneyness-maturity categories since we have lower maturities and therefore not as large levels of moneyness available. The maturities

[11] Doing the same tests on a larger set of data (the number of options used for the regressions is 1090 for S&P500 and 1209 for Apple) would help to conclude on the significance of β_4.

are divided as: DTM$\leqslant$20, 20<DTM$\leqslant$40 and DTM>40. For moneyness, the classification is: (1) DOTM (M<0.99 for S&P500 and M<0.97 for Apple); (2) OTM (0.99<M<1 for S&P500 and 0.97<M<1 for Apple); (3) ITM (1<M<1.01 for S&P500 and 1<M<1.03 for Apple); (4) DITM (M>1.01 for S&P500 and M>1.03 for Apple).

Table H.1 in appendix details the number of options used. The results are displayed on tables H.2 and H.3 in appendix and are compared with the results without re-calibration.

S&P500 The main observation is that Heston and HN are still systematically beaten by BS and PBS. Also, quite surprisingly, the errors are nearly always larger with re-calibration of parameters than without. This is counter-intuitive since one could think, at least for HN and Heston, that frequently re-estimating the parameters would have counter-balanced the over-fitting issue that made these models perform badly for hedging. In the contrary, our results indicate that time-varying parameters do not generate more accurate Delta estimates. To improve HN and Heston hedging performance, we believe that the most promising solution would be to use a hedging-based estimation function rather a pricing-based one.

Apple The effect of re-calibration is quite different than for the S&P500. For OTM, ITM and DITM options, we do not observe any significant differences depending on whether the parameters are frequently re-calibrated or not. However, for DOTM options, while with constant parameters HN and Heston are outperformed by BS and PBS, this is interestingly not the case anymore with re-calibration: the ranking becomes HN, Heston, PBS and lastly BS, for all maturities. Regarding the amplitude of errors, we now have that re-calibrating parameters is much more worthwhile than for the S&P500. Lower errors are frequently observed, and especially for DOTM options: 8 out of 12 categories generate more accurate hedging than with constant parameters.

Overall, we can thus conclude that frequently re-estimating parameters throughout the life of the option is worthwhile only in the case of sufficiently unstable underlying asset returns (Apple), and especially for DOTM options. In that case, Heston and HN outperform BS and PBS for DOTM options as well. If the underlying asset volatility is low (S&P500), then one is better off by not including variations in the parameters.

6.4 The Over-Hedging Behaviour: Explanation from Simulated Data

This final section answers to our third research sub-question: *How does under/over-hedging depend on the re-balancing frequency, the underlying asset expected return and the underlying asset volatility?*.

We previously observed a systematic over-hedging on our data, that we now wish to further understand. To do so, we follow the methodology developed in section 5.4 and perform Monte Carlo simulations on HN. This enables to have a lot of flexibility on the variables whose effects we want to understand. Understanding when and why an over-hedging appears is relevant since having an over-hedged portfolio guarantees that you do not fall short of the payoff at maturity, while even leaving you with a positive bonus that will help to cover the transaction costs that can not be ignored in practice. Conversely, an under-hedged portfolio is not desirable for practitioners.

The most plausible explanation for the observed over-hedging effect results from the fact that the price of a call option (and incidentally also of a put option) is a convex function of the stock price (Joshi, 2008), i.e. the Gamma ($\Gamma_t \equiv \frac{\partial \Delta_t}{\partial S_t}$) is always positive. This means that Δ_t, which is the slope of the call option function, will increase more than it will decrease for the same absolute percentage change in the stock price. This has a positive effect on the replicating portfolio value. This is pointed out by Joshi (2008), who explains that being long Gamma means that the procedure of hedging will make money during the life of the option. If we were able to apply continuous re-balancing, then this would have no effect. However, we only re-balance the portfolio on a daily basis. Since our procedure accounts for the whole life of the option, the convexity effects have the time to accumulate and therefore we very often over-hedge the payoff at maturity, especially for a long maturity. This would also explain why the literature rather points to an under-hedging effect (as explained at the beginning of section 6.3.1): the short hedging horizon makes that convexity has a more minor effect.

In addition, we argue that a higher volatility of the underlying asset should lead to more over-hedging since the returns are larger in absolute value, which increases the convexity effect. However, we saw previously that Apple did not always lead to more over-hedging than the S&P500. This may be explained by the fact, as figure 5.1 shows, that the S&P500 had increased in value over the time period considered while Apple had decreased. This should *ceteris paribus* lead to a larger over-hedging for S&P500 options because a lower expected return means that there is less effect coming from magnified positive daily returns by the positive Gamma.

In brief, we wish to test the effect of the re-balancing frequency, expected return and underlying volatility on under/over-hedging.

Table I.1 in appendix reports the normalized \$ME for our different levels of re-balancing frequency, expected return and volatility. The strongest result is that the normalized \$ME always increases when the replicating portfolio is re-balanced less frequently, linked to the positive Gamma effect. When re-balancing is done twice a day, we always observe an under-hedging, while a normalized \$ME very close to 1 is obtained when re-balancing every five days. Regarding volatility, when going from $\sigma = 15\%$ to 30%, the errors systematically increase for a twice a day re-balancing, and decrease otherwise. From 30% to 45% volatility, the errors overall tend to increase. The hypothesis that a larger volatility increases the proportion of over-hedging is thus not strongly confirmed. Finally, the effect of the expected return depends on the re-balancing frequency. When going from $\mu = -10\%$ to 10%, the errors overall tend to decrease for a frequent re-balancing (twice a day and daily) and tend to increase otherwise. Our hypothesis that a higher expected return leads to a higher proportion of over-hedging is therefore only verified when re-balancing is not too frequent.

Figure 6.9 illustrates the the payoff replicating strategy when $\sigma = 30\%$ and $\mu = 5\%$ for our four different re-balancing frequencies. It shows that the the payoff at maturity is over-hedged under a low frequency, and the contrary for a high frequency. Interestingly, the distance between the payoff and replicating portfolio value at maturity gets lower when the frequency increases.[12] This can simply be explained by the fact that, under a lower re-balancing frequency, larger changes in the underlying stock price can occur between two revision dates. In turn, this can substantially change the value of the Delta

[12]More precisely, in our case with $\sigma = 30\%$ and $\mu = 5\%$, the payoff replication \$MAE takes the following values from the highest frequency (twice a day) to the lowest (every five days) under 1000 Monte Carlo random paths: 2.4672, 2.5407, 3.1936 and 6.7544.

FIGURE 6.9: Payoff replication strategy for at-the-money options simulated under the HN model for different re-balancing frequencies.

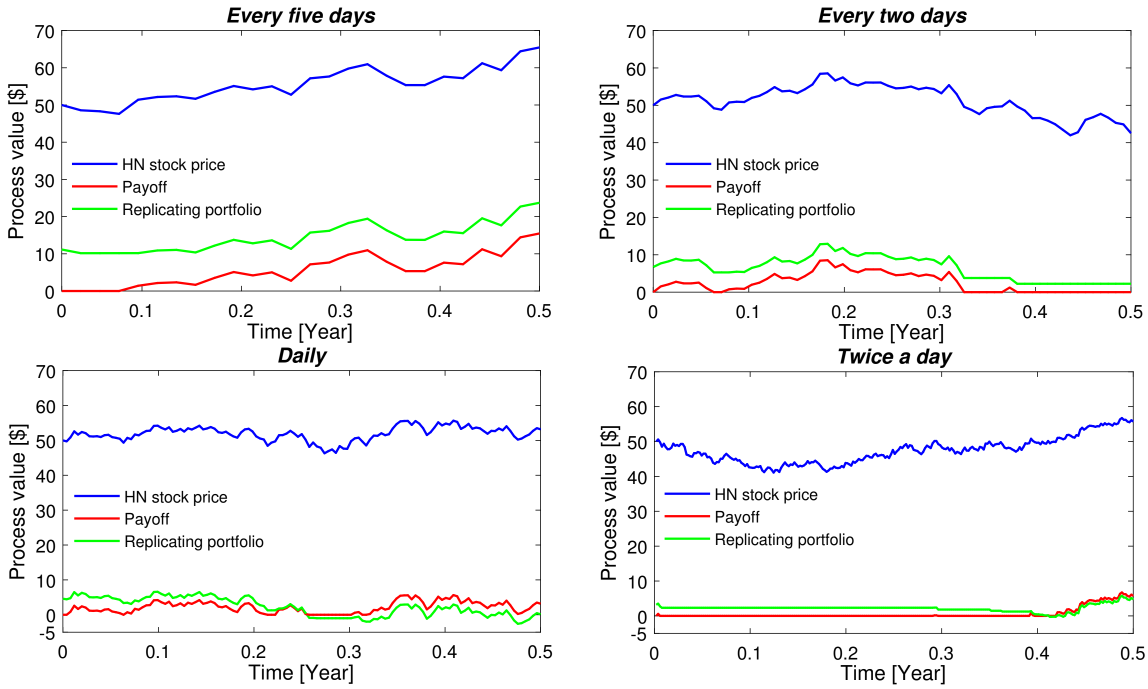

Notes: The simulations were performed with $S_0 = K = 50$, $T = 0.5$, $\mu = 5\%$ and $\sigma = 30\%$. The mean and volatility are reflected by the HN parameters' values as follows: λ is set such that $\mu = r + 252 * k * \lambda * \bar{h}$, β_2 such that $\sigma = \sqrt{252 * k * \bar{h}}$ where $k = 0.5, 1, 2, 5$ depending on the re-balancing frequency, $\bar{h} = \beta_0/(1 - \beta_1 - \beta_2\gamma^2)$, i.e. the long run variance under the HN model, and $h_0 = \bar{h}$. The other parameters are: $r = 3\%$, $\beta_0 = 8 * 10^{-6}$, $\beta_1 = 0.45$, $\gamma = 700$.

while the number of stocks held in the portfolio is not adjusted accordingly between the two re-balancing dates. This ultimately makes hedging less effective.

Chapter 7

Conclusion

This thesis wished to provide a thorough study of the hedging performance and behaviour, through payoff replication, of four well-known option pricing models - Black-Scholes, Practitioner Black-Scholes, Heston-Nandi and Heston - each of which distinguishes by the way the volatility of the underlying asset is modeled. The results were obtained on two time series (one per volatility category): the S&P500 index and Apple stock, based on option quotes from January 4 to March 30, 2016.

The introduction of the thesis presented the research question that we wanted to tackle. In short, we wanted to assess the relative performance of our four models regarding out-of-sample payoff replication, as well as how the moneyness, maturity and volatility of the underlying impact the performance. Additionally, three sub-questions appeared as very relevant from our study: whether an accurate fit of the market option prices is a good indicator of the hedging performance, the effect of a frequent re-calibration of parameters and the drivers of under/over-hedging. After reviewing the relevant theoretical content in chapters 2, 3 and 4, chapter 5 detailed the methodology followed and chapter 6 finished with the presentation and discussion of the results.

Chapter 1 already summarized our main results that chapter 6 detailed and discussed them extensively, also in light of the literature. Therefore, this conclusion will focus on the implications of the results, and will end up with the limitations of our work as well as suggestions for future research.

7.1 Implications of the Results

From the payoff replication investigation performed in this thesis, we wish to draw the reader's attention to three implications.

The first implication arising from our research is that there is not one ideal model: the best model to use depends on the problem at hand. Specifically, it depends on the purpose of the exercise (e.g. pricing or hedging), the specificities of the underlying (and especially its volatility), the maturity of the option, its moneyness and finally the procedure followed (e.g. whether the parameters are frequently re-calibrated or how the parameters are estimated). Several examples erupt from our results, such as:

- Under a high underlying volatility, when the parameters are kept constant throughout the life of the option, Heston and HN are outperformed by BS and PBS for DOTM options. In the contrary, when applying a frequent re-calibration of parameters, Heston and HN beat BS and PBS for those same options.
- Under a low volatility, HN and Heston are always outperformed by BS and PBS. Conversely, under a higher volatility and for options sufficiently OTM, Heston and HN perform best for long-term contracts.
- For pricing purposes, HN and Heston perform overall better than PBS and BS. However, under a low underlying volatility, PBS sometimes beats them for DITM and DOTM options. Under a high volatility, PBS does better for long-term options as well as DITM and DOTM mid-term ones.

Therefore, practitioners have to be careful not to blindly follow only one model for all purposes. They must carefully perform historical analyses to understand which model they should use for which underlying asset, which type of option and which purpose.

The second implication is that adding complexity to make the procedure and the models more realistic may counter-intuitively lead to worse results. Indeed, our findings support this implication. On the one hand, we have shown that more realistic volatility specifications tend to decrease the hedging accuracy due to an over-fitting issue. On the other hand, there is some evidence from our results that a frequent re-calibration of parameters, although more realistic since it accounts for more updated information, sometimes leads to larger payoff replication errors, and even nearly systematically under the S&P500. Therefore, in light of such counter-intuitive results, practitioners have to strive to rely as much as possible on analytical reasoning rather than common sense. As for the the first implication, well-designed historical tests are necessary to that extent.

The third and most important implication is the importance of application-oriented calibration procedures. Indeed, our results have shown that HN and Heston, which are very good performers for out-of-sample pricing, display quite poor results for out-of-sample payoff replication. The reason is that the good pricing performance comes from an over-fit of option prices that in turn generates poor Delta estimates. Therefore, we argue that using a pricing estimation loss function, which is the standard practice, is at odds with the hedging exercise for which the options are applied. This was pointed out by Christoffersen and Jacobs (2003): "We argue that the relevant discussion should not start out by discussing how to estimate the parameters, but rather by stating what the evaluation (typically out-of-sample) loss function is. This loss function is dictated by the purpose of the empirical exercise and can be related to a hedging problem or a risky investment strategy for example." Therefore, the focus for practitioners should not altogether be put on which model to use but also on how to build a calibration procedure that fits the final purpose of the practitioner. In other words, we point to the importance of a consistency between the estimation and evaluation loss functions. In our case, a hedging-based estimation loss function should therefore ideally be used.

7.2 Limitations of the Research

This section points out the three main limitations of the research that we performed, and that we want the reader to be aware of.

The first limitation arises from the fact that we ignore transaction costs when re-balancing the Delta-neutral portfolio, similarly to what is done in the literature. However, in practice, they can not be ignored. Therefore, one should ideally test whether our findings are robust to including transaction costs in the replicating portfolio process. Depending on the size of these costs, a less frequent re-balancing may prove more viable, at the risk of having larger stock price changes between two revision dates. Depending on how the Delta reacts to changes in the underlying (i.e. the value of Gamma) under each model, we may have some models generating less costs than others, hence making them relatively more attractive than when ignoring such costs.

The second limitation comes from the fact that our conclusions are based on only two time series, one per volatility category, and whose option quotes cover a small time span of three months. This period, as mentioned in section 5.1, is substantially smaller than most of the studies in the literature. This originates from the availability of historical option quotes on Bloomberg which is limited to three months. To mitigate this, we collected the data on a daily frequency rather than weekly in the literature. Still, the reader should keep in mind that the results and conclusions from our work may not always be robust to a larger time period, and more importantly to other underlying asset samples. Ideally, one should verify whether our results are verified on a longer time span and on other time series having different statistical behaviours. However, we can reasonably expect that our general conclusions and implications should not drastically be changed but that it may help to refine them more precisely and potentially to find additional conclusions.

The third limitation comes from our data collection. By collecting closing prices for both the options and the underlying asset, we may face a non-synchronous bias, which arises when the option quote and the corresponding underlying price are not recorded simultaneously. The CBOE, where our options are traded, closes at 3:15 p.m., while both the NASDAQ (where Apple is traded) and the S&P500 stop trading at 4:00 p.m. However, as mentioned in section 5.1, we mitigated this bias by taking a large enough number of daily observations. Also, Bates (2000), in a decomposition of option pricing residuals into measurement and specification error based on S&P500 futures options, found that the latter is from one to four times larger than the former. This makes us confident that the non-synchronous bias should not seriously affect our results.

7.3 Suggestions for Future Research

Finally, we end up this work with three suggestions for future research that we believe are worth investigating.

First, other models than those to which we restricted in this thesis could be analyzed. One could consider other functional forms for the PBS model for example. More importantly, we consider that jump models, which include random discontinuous changes in the underlying price dynamics, are a path worth following. Indeed, Bates (2003) pointed out that "to simultaneously summarize volatility evolution and capture the fat-tailed properties of daily returns, [...] jumps in returns and/or in volatility, and probably both" are needed. It would then be interesting to assess whether jump models are able to generate an accurate hedging performance and how this model would behave regarding payoff replication. As mentioned in section 4.3.5, Bakshi et al. (1997) found that adding jumps on top of stochastic volatility does not improve the hedging performance. However, our empirical results showed that we do not always agree with the literature findings, so that it is still worth applying our procedure to a jump model in the context of our research questions. For example, one could opt for the first jump model developed by Merton (1976) where the underlying stock returns are a mixture of continuous and jump processes. This model is attractive model because it generates a closed-from formula for a European option price.

A second suggestion is to consider other hedging strategies than only the simple Delta-neutral one. Indeed, while theoretically this simple strategy is enough to eliminate all risks in the BS framework, this is not so for other models such as HN and Heston which are market incomplete since volatility risk is also present. Of course, the possibility to eliminate all risks with a simple Delta-neutral strategy is even less true empirically. Therefore, it would be interesting to compute the payoff replication errors when other Greeks than only Delta are hedged. Two common strategies to that extent are to construct Delta-Gamma and Delta-Vega neutral portfolios. A Delta-Gamma neutral portfolio is hedged both with respect to a change in the stock price, but also to a change in the value of Delta, meaning that "Gamma neutrality provides protection against larger movements in the stock price between hedge re-balancing" (Hull, 2012). To achieve this, it is necessary to trade one additional option. If this option has a Gamma equal to Γ and the Delta-neutral portfolio's Gamma is Γ' then, if we add a quantity w of the option, the portfolio will have a Gamma equal to $w\Gamma + \Gamma'$, which makes that w must be worth $-\Gamma'/\Gamma$ to make the portfolio Gamma neutral (Hull, 2012). This Delta-Gamma neutral portfolio must be periodically re-balanced to keep a Gamma of 0 as the values of Γ and Γ' evolve. Racicot and Théoret (2006) showed, in the context of the BS model, that Delta-Gamma hedging is "quite superior to a simple Delta hedging." Along the same lines, one can construct a Delta-Vega neutral portfolio to hedge against the changes in the underlying asset volatility. Bakshi et al. (1997) showed that Delta-Vega hedging produces far better hedging results, and that it makes stochastic volatility models outperform BS for ITM options. This may counter the fact that we always observed that HN and Heston were outperformed by BS and PBS for those options. Yung and Zhang (2003) also found that Delta-Vega hedging increases the relative performance of the EGARCH model compared to BS. These results show that more advanced hedging procedures are worth investigating.

The third and final suggestion for future research is linked to the third implication mentioned before on the importance of having consistent estimation and evaluation loss functions. It is worth assessing how the results are changed when calibrating parameters following a hedging-based estimation loss function, for example based on payoff replication. We believe that the results may be quite seriously different under such a procedure compared to ours. That being said, computational constraints also have to be considered and may make such loss functions hard to implement in practice. Indeed, one significant advantage of the models analyzed in this thesis is that closed-form formulas are available for option prices, which makes the estimation procedure based on pricing loss functions very computationally-friendly. Conversely, using a loss function based on payoff replication, for example, would lead to a substantially larger computational burden since payoff replication is a path-dependent criterion, i.e. it depends on the path taken by the underlying asset price until expiration of the option. This path-dependency also implies that the estimation would need to be made on option data quite far in the past rather than on a few days ago, which is also not so attractive. Therefore, one would have to balance the improvement in hedging accuracy with the increasing complexity of the implementation procedure.

Bibliography

Bakshi, G., Cao, C., & Chen, Z. (1997). Empirical Performance of Alternative Option Pricing Models. *Journal of Finance*, 52, 2003-2049.

Bates, D. (2000). Post-'87 Crash Fears in the S&P 500 Futures Option Market. *Journal of Econometrics*, 94, 181-238.

Bates, D. (2003). Empirical Option Pricing: A Retrospection. *Journal of Econometrics*, 116, 387-404.

Berkowitz, J. (2002). Forecasting Option Values with Unobservable Volatility and False Models. Working Paper, University of Houston.

Black, F. (1976). *Studies of Stock Price Volatility Changes.* Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section, 177–181.

Black, F., & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *The Journal of Political Economy*, 81(3), 637–654.

Board of Governors of the Federal Reserve System. (2016). *Historical data.* Retrieved from `http://www.federalreserve.gov/releases/h15/data.htm`.

Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 31, 307–327.

Bollerslev, T., Ray, C., & Kenneth, K. (1992). ARCH Modeling in Finance : A Review of the Theory and Empirical Evidence. *Journal of Econometrics*, 52, 5–59.

Breusch, T. (1978). Testing for Autocorrelation in Dynamic Linear Models. *Australian Economic Papers*, 17, 334–355.

Breusch, T., & Pagan, A. (1979). A Simple Test for Heteroscedasticity and Random Coefficient Variation. *Econometrica*, 47(5), 1287–1294.

Brigo, D., & Mercurio, F. (2006). *Interest Rate Models - Theory and Practice* (2nd ed.). Berlin: Springer.

CBOE. (2016). *S&P500 Index Options (SPX)*. Retrieved from http://www.cboe.com/framed/pdfframed.aspx?content=/micro/spx/pdf/spx_qrg2.pdf&section=SECT_MINI_SITE&title=SPX+Fact+Sheet.

Chance, D. (2003). *Analysis of Derivatives for the CFA Program*. Baltimore: AIMR.

Christoffersen, P. (2012). *Elements of Financial Risk Management* (2nd ed.). Oxford: Academic Press.

Christoffersen, P., & Jacobs, K. (2003). The Importance of the Loss Function in Option Valuation. *Journal of Financial Economics*, 72, 291–318.

Christoffersen, P., & Jacobs, K. (2004). Which GARCH Model for Option Valuation? *Management Science*, 50(9), 1204–1221.

Cox, J., Ingersoll, J., & Ross, S. (1985). A Theory of the Term Structure of Interest Rates. *Econometrica*, 3, 145–166.

Cox, J., & Ross, S. (1976). The Valuation of Options for Alternative Stochastic Processes. *Journal of Financial Economics*, 3, 145–166.

Crisostomo, R. (2014). An Analysis of the Heston Stochastic Volatility Model: Implementation and Calibration using Matlab. *CNMV Working Paper*, 58.

Dennis, P., & Mayhew, S. (2000). Implied Volatility Smiles: Evidence from Options on Individual Equities. Working Paper, University of Virginia.

Duan, J. (1995). The GARCH Option Pricing Model. *Mathematical Finance*, 5(1), 13–32.

Dumas, B., Fleming, J., & Whaley, R. (1998). Implied volatility functions: Empirical tests. *Journal of Finance*, 53, 2059–2106.

Engle, R., & Patton, A. (2001). What Good is a Volatility Model? *Quantitative Finance*, 1(2), 237–245.

Feller, W. (1951). Two Singular Diffusion Problems. *Annals of Mathematics*, 54, 173–182.

FIA. (2015). *2014 FIA Annual Global Futures and Options Volume: Gains in North America and Europe Offset Declines in Asia-Pacific*. Retrieved from https://fimag.fia.org/articles/2014-fia-annual-global-futures-and-options-volume-gains-north-america-and-europe-offset.

Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide*. Hoboken: John Wiley & Sons.

Godfrey, L. G. (1978). Testing Against General Autoregressive and Moving Average Error Models when the Regressors Include Lagged Dependent Variables. *Econometrica*, 46(6), 1293–1301.

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

Heston, S., & Nandi, S. (2000). A Closed-Form GARCH Option Valuation Model. *The Review of Financial Studies*, 13(3), 585–625.

Hsieh, K., & Ritchken, P. (2005). An Empirical Comparison of GARCH Option Pricing Models. *Review of Derivatives Research*, 8(3), 129–150.

Hull, J. (2012). *Options, Futures and Other Derivatives* (8th ed.). Boston: Pearson Prentice Hall.

Hull, J., & White, A. (1987). The Pricing of Options on Assets with Stochastic Volatilities. *The Journal of Finance*, 42(2), 281–300.

Investopedia (n.a.). *Treasury Bill - T-Bill*. Retrieved from http://www.investopedia.com/terms/t/treasurybill.asp.

Jackel, P. (2004). *Stochastic Volatility Models: Past, Present and Future*. Retrieved from http://www.pjaeckel.webspace.virginmedia.com/StochasticVolatilityModels-PastPresentAndFuture.pdf

Jarque, C., & Bera, A. (1987). A Test for Normality of Observations and Regression Residuals. *International Statistical Review*, 55(2), 163-172.

Joshi, M. (2008). *The Concepts and Practice of Mathematical Finance* (2nd ed.). Cambridge: Cambridge University Press.

Kermiche, L. (2014). Too Much of a Good Thing? A Review of Volatility Extensions in Black-Scholes. *The Journal of Applied Business Research*, 30(4), 1171–1182.

Lehar, A., Scheicher, M., & Schittenkopf, C. (2002). GARCH vs Stochastic Volatility: Option Pricing and Risk Management. *Journal of Banking and Finance*, 26(1), 323–345.

Liu, J., & Pan, J. (2003). Dynamic Derivative Strategies. *Journal of Financial Economics*, 69, 410–430.

Macbeth, J., & Merville, L. (1979). An Empirical Examination of the Black-Scholes Call Option Pricing Model. *Journal of Finance*, 5, 1173–1186.

Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *The Journal of Business*, 36(4), 394–419.

Merton, R. (1976). Option Pricing when Underlying Stock Returns are Discontinuous

Nelson, D. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica*, 59(2), 347–370.

Newey, W. K., & West, K. D. (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3), 703–708.

Ortega, J. (2009). *GARCH Option Pricing via Local Risk Minimization.* Retrieved from http://cermics.enpc.fr/cnf/Ortega.pdf.

Petitjean, M., & Moyaert, T. (2011). The Performance of Popular Stochastic Volatility Option Pricing Models During the Subprime Crisis. *Applied Financial Economics*, 21(14), 1059–1068.

Racicot, F., & Théoret, R. (2006). Simulations de la couverture delta et de la couverture delta-gamma d'un portefeuille dans le cadre du modèle de Black et Scholes. *RePAd Working Paper Series*, Département des sciences administratives, UQO.

Rosu, I., & Stroock, D. (2004). On the Derivation of the Black-Scholes Formula. *Séminaire de Probabilités*, 37, 399–414.

Rouah, F. (n.a.). *Euler and Milstein Discretization.* Retrieved from http://frouah.com/pages/finmath.html.

Rouah, F. D., & Vainberg, G. (2007). *Option Pricing Models and Volatility Using Excel-VBA*. Hoboken: John Wiley & Sons.

S&P Dow Jones Indices. (2016a). *S&P500.* Retrieved from http://us.spindices.com/indices/equity/sp-500.

S&P Dow Jones Indices (2016b). *S&P500 Dividend Points Index.* Retrieved from http://us.spindices.com/indices/equity/sp-500-dividend-points-index.

S&P Global. (2016). *S&P U.S. Indices Methodology.* Retrieved from https://us.spindices.com/documents/methodologies/methodology-sp-us-indices.pdf.

Stein, E., & Stein, J. (1991). Stock Price Distributions with Stochastic Volatility: An Analytic Approach. *The Review of Financial Studies*, 4(4), 727–752.

van der Wijst, N. (2013). *Finance: A Quantitative Introduction.* Cambridge: Cambridge University Press.

Vrins, F. (2015). *Fixed Income Analysis (ii): Options and Callable Products.* Louvain School of Management, Course Material.

Yung, H., & Zhang, H. (2003). An Empirical Investigation of the GARCH Option Pricing Model: Hedging Performance. *The Journal of Futures Markets*, 23(12), 1191–1207.