Numerical and Empirical Studies of Option Pricing

A thesis submitted to The University of Manchester for the degree of Doctor of Philosophy in the Faculty of Humanities.

2014

Przemyslaw Stanislaw Stilger

Manchester Business School

Contents

Word count: 33,846

List of Tables

List of Figures

The University of Manchester

Przemyslaw Stanislaw Stilger

Doctor of Philosophy

20 October 2014

ABSTRACT

This thesis makes a number of contributions in the derivative pricing and risk management literature and to the growing literature that exploits information embedded in option prices. First, it develops an effective numerical scheme for importance sampling scheme of Fouque and Tullie (2002) based on a 2-dimensional lookup table of stock price and time to maturity that dramatically improves the speed of this importance sampling scheme. Second, the thesis presents an application of this importance sampling scheme in a Multi-Level Monte Carlo simulation. Such combination yields greater variance reduction compared to Multi-Level Monte Carlo or importance sampling alone. Third, it demonstrates how the Greeks can be computed using the Likelihood Ratio Method based on characteristic function, and how combining it with importance sampling leads to a significant variance reduction for the Greeks. Finally, it documents the positive relationship between the risk-neutral skewness (RNS) and future realized stock returns that is driven by the underperformance of highly negative RNS portfolio. The results provide strong evidence that the underperformance of stocks with the lowest RNS is driven by those stocks that are associated with a higher hedging demand, relative overvaluation and are also too costly or too risky to sell short. Moreover, by decomposing RNS into its systematic and idiosyncratic components, this thesis shows that the latter drives the positive relationship with future realized stock returns.

Declaration

No portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright

Acknowledgment

Many thanks to my supervisors Prof. Ser-Huang Poon and Dr. Alexandros Kostakis for their guidance and support. I would also like to thank Dr. Simon Acomb and Dr. Eberhard Mayerhofer for helpful discussions. I thank the Economic and Social Research Council [ES/J500094/1] as well as Manchester Business School for their financial support. Finally, I would like to thank my family for their love and support.

Chapter 1

Introduction

This thesis concerns the numerical and empirical aspects of derivative pricing. The first part of the thesis focuses on improving the efficiency of Monte Carlo estimators through importance sampling and multi-level simulations. We develop an effective numerical scheme for importance sampling scheme of Fouque and Tullie (2002) transforming out-of-the-money sample paths into in-the-money ones thereby improving the derivative price estimation. We have also extended this importance sampling scheme so that it can accommodate more complex and realistic price dynamics. We have further improved the speed of computation by incorporating Multi-Level Monte Carlo estimators. We demonstrated the usefulness of our work by implementing some state-of-the-art financial models such as Heston's stochastic volatility, Bates's stochastic volatility with jumps, Heston-Hull-White and Heston-Cox-Ingersoll-Ross with stochastic volatility and stochastic interest rates. We demonstrated the power of the enhanced importance sampling scheme by pricing single asset options such as European and barrier options as well as multi-asset options such as basket, best-of, worst-of, spread, absolute, composite and quotient options. In order to estimate the Greeks, we introduce the Likelihood Ratio Method based on characteristic function so that we can achieve a significant variance reduction for the Greeks by incorporating the importance sampling scheme developed here.

The second part of this thesis concerns the relation between the risk-neutral skewness (RNS) of the returns distribution and future realized stock returns. We use daily option prices for a large sample of US stocks during the period 1996-2012 and estimate RNS of each stock using the model-free methodology of Bakshi, Kapadia and Madan (2003). We find significant and robust evidence that RNS is positively related to future realized stock returns. A strategy that is long the quintile portfolio with the highest RNS stocks and short the quintile portfolio with the lowest RNS stocks, rebalanced monthly, yields an average return of 61 basis points (t-stat: 2.25) per month, Fama-French (FF) alpha of 49 basis points (t-stat: 2.23) per month and Fama-French-Carhart (FFC) alpha of 55 basis points (t-stat: 2.52) per month. The return of this strategy is mainly driven by the underperformance of highly negative RNS portfolio. We further decompose RNS into its systematic and idiosyncratic components and find that the latter drives the positive relationship between RNS and future realized stock returns. A strategy that is long the quintile portfolio with the highest risk-neutral idiosyncratic skewness (RNIS) stocks and short the quintile portfolio with the lowest RNIS stocks, rebalanced monthly, yields an average return of 62 basis points (t-stat: 2.40) per month, Fama-French (FF) alpha of 51 basis points (t-stat: 2.39) per month and Fama-French-Carhart (FFC) alpha of 57 basis points (t-stat: 2.68) per month. We find robust evidence that the significant underperformance of stocks exhibiting the most negative RNS is driven by those stocks that are relatively overpriced, too costly or too risky to sell short and characterized by high hedging demand.

This thesis makes a number of contributions in the derivative pricing and risk management literature and to the growing literature that exploits information embedded in option prices.

First, we propose an effective numerical scheme for importance sampling scheme of Fouque and Tullie (2002) which requires Black-Scholes option price and delta at every time step for different pairs of stock price and time to maturity. Our approach relies on a 2-dimensional lookup table of stock price and time to maturity that significantly reduces the number of times Black-Scholes option prices and deltas are calculated thereby dramatically improving the speed of importance sampling. Moreover, we have extended this importance sampling scheme to accommodate more complex and realistic price dynamics such as Heston's stochastic volatility, Bates's stochastic volatility with jumps, Heston-Hull-White and Heston-Cox-Ingersoll-Ross with stochastic volatility and stochastic interest rates. We have also successfully applied this extended importance sampling scheme to price single asset options such as European and barrier options as well as multi-asset options such as basket, best-of, worst-of, spread, absolute, composite and quotient options.

Second, we apply the importance sampling scheme of Fouque and Tullie (2002) in a Multi-Level Monte Carlo simulation. Here, the improvement comes from lower computational cost of Multi-Level Monte Carlo compared to the basic Monte Carlo and variance reduction through importance sampling.

Third, we show how the Greeks can be computed using the Likelihood Ratio Method based on characteristic function, and how combining it with importance sampling leads to a significant variance reduction for the Greeks.

Finally, we document the positive relationship between RNS and future realized stock returns that is driven by the underperformance of highly negative RNS portfolio. We demonstrate that the underperformance of stocks with the lowest RNS is driven by those stocks that are simultaneously characterized by higher hedging demand, relative overvaluation and that are also too costly or too risky to sell short. Moreover, we further decompose RNS into its systematic and idiosyncratic components and find the latter responsible for driving the positive relationship between RNS and future realized stock returns.

The rest of this thesis is organized as follows. Chapter 2 provides an overview of the Monte Carlo method in the context of derivative pricing and variance reduction techniques. We start with a definition of the Monte Carlo estimator and its statistical properties such as unbiasedness and consistency. This is followed by a discussion of a few variance reduction techniques such as importance sampling, control variate, antithetic variates, conditional Monte Carlo, stratified sampling and Multi-Level Monte Carlo.

In Chapter 3, we apply importance sampling to Heston's stochastic volatility model and Bates's stochastic volatility model with jumps using European call and down-and-out put options as examples. Moreover, we propose an effective numerical scheme that dramatically improves the speed of importance sampling. Finally, we show how the Greeks can be computed using the Likelihood Ratio Method, and how combining it with importance sampling leads to a significant variance reduction for the Greeks.

In Chapter 4, we present an application of importance sampling to multi-asset options under the Heston and the Bates models as well as to the Heston-Hull-White and the Heston-Cox-Ingersoll-Ross models. Moreover, we provide an efficient importance sampling scheme in a Multi-Level Monte Carlo simulation. In all cases, we explain how the Greeks can be computed in the different simulation schemes using the Likelihood Ratio Method, and how combining it with importance sampling leads to a significant variance reduction for the Greeks.

In Chapter 5, we study the relationship between RNS and future realized stock returns. We find significant and robust evidence that RNS is positively related to future realized stock returns. Decomposing RNS into its systematic and idiosyncratic components, we find that the latter drives the positive relationship with future stock returns. To explain the positive relationship between RNS and future realized stock returns, we put forward a mechanism on the basis that certain stocks are relatively overpriced and too costly or too risky to sell (short-sell). As a result, hedgers and speculators resort to the options market as means of hedging their underlying positions and/or trading on their negative expectations, driving RNS to very low negative values. In this way, option prices contain information that it is not embedded in stock prices. As this information is diffused to the stock market over time, relatively overpriced stocks subsequently yield negative returns, giving rise to a positive relationship between RNS and future realized stock returns.

Finally, Chapter 6 concludes with a summary of the main findings and contributions of this thesis.

Chapter 2

Monte Carlo Methods for Derivative Pricing

2.1 Monte Carlo

Derivatives are financial contracts whose value is derived from an underlying instrument such as stock, bond, commodity, interest rate, currency, loan or index. For example, a European call (put) option gives the holder the right to buy (sell) the underlying for a fixed price K at time T. Commonly used techniques for pricing derivatives include: analytical solution, transform methods, numerical integration, finite differences, multinomial lattices or Monte Carlo method. The flexibility of the Monte Carlo method and its ease of implementation make it powerful and popular in both derivatives pricing and risk management. To obtain a Monte Carlo estimate of a derivative price, one simulates n sample paths of the stock price process S under the risk-neutral probability measure Q and then computes the payoff of the derivative for each realization. Finally, the derivative price is calculated as the discounted arithmetic average of the payoff associated with each sample path. The convergence of the Monte Carlo estimate to the true derivative price is guaranteed by the Law of Large Numbers and the rate of convergence $\sqrt{n}$ is given by the Central Limit Theorem. The Monte Carlo estimator is unbiased as its expectation equals the true derivative price. It is also consistent, because it converges in probability to the true derivative price. By the probabilistic nature of the Monte Carlo, a precise estimate of the derivatives price requires a large number of sample paths. The quality of the Monte Carlo estimate is usually quantified by a confidence interval. The objective of the first part of this thesis is to improve the accuracy and efficiency of the Monte Carlo estimator for state-of-the-art derivative pricing models.

Let the stock price process be given by the stochastic differential equation (SDE) below

$$dS_t = a(t, S_t)\, dt + b(t, S_t)\, dW_t,\ t \in [t_0, T] \tag{2.1}$$

where dW_t is a Brownian motion. Functions a and b are the drift and the diffusion terms respectively. Note that both a and b are functions of stock price and time. For the time interval $[t_0, T]$, we fix a time discretization of the form

$$0 = t_0 < t_1 < \cdots < t_m = T$$

where m is the number of time steps. We denote the discrete approximation of the stock price in (2.1) at time t_j by $S(t_j)$.

Two of the most popular discretization schemes are the Euler scheme and the Milstein scheme. The Euler scheme is given by

$$S(t_{j+1}) = S(t_j) + a(t_j, S(t_j))\, \triangle t + b(t_j, S(t_j))\, \triangle W$$

where $\triangle t = t_{j+1} - t_j$ and $\triangle W = W_{j+1} - W_j$ for $j = 0, 1, \cdots, m-1$. The Euler scheme is the simplest discretization scheme for (2.1) which expands the drift term to $O(\triangle t)$, but the diffusion term to only $O\left(\sqrt{\triangle t}\right)$. The Milstein scheme expands both the drift term and diffusion term to $O(\triangle t)$ and it is given by

$$\begin{aligned} S(t_{j+1}) &= S(t_j) + a(t_j, S(t_j))\, \triangle t + b(t_j, S(t_j))\, \triangle W \\ &+ \frac{1}{2} b'(t_j, S(t_j))\, b(t_j, S(t_j)) \left[\triangle W^2 - \triangle t\right] \end{aligned}$$

where $b'(t, S) = \frac{\partial}{\partial s} b(t, S)$.

The first type of error of the Monte Carlo estimator is the discretization error caused by the numerical integration of the SDE which tends to zero as the number of time steps increases, i.e. as $m \to \infty$. The second type of error is the statistical error embedded in the random draws of the simulation. According to the Central Limit Theorem, as the number of draws increases, i.e. as $n \to \infty$, we have convergence in distribution

$$\sqrt{n} \left(\frac{\hat{P} - P}{\sigma} \right) \rightsquigarrow N(0, 1)$$

where n is the total number of sample paths, P is the true derivative price, $\hat{P}$ is its Monte Carlo estimate given by $\frac{1}{n} \sum_{i=1}^{n} P_i$, σ is the sample standard deviation given by $\sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(P_i-\hat{P}\right)^2}$ and $N(0,1)$ denotes the standard normal distribution. It follows that

$$\left[\hat{P}-1.96\frac{\sigma}{\sqrt{n}},\ \hat{P}+1.96\frac{\sigma}{\sqrt{n}}\right]$$

is the 95% confidence interval of the Monte Carlo estimator. We note that the length of the confidence interval is proportional to σ and inversely proportional to $\sqrt{n}$. The ratio $\frac{\sigma}{\sqrt{n}}$ is the standard error of the Monte Carlo estimator. It decreases with $\sqrt{n}$ regardless of the number of dimensions, which implies that the Monte Carlo estimator can achieve any desired level of accuracy and the confidence interval can be made arbitrarily small by increasing n.

There are various variance reduction techniques for reducing the variance of the Monte Carlo estimator and improving its efficiency. In the following sections, we briefly discuss a few variance reduction techniques such as importance sampling, control variate, antithetic variates, conditional Monte Carlo, stratified sampling and Multi-Level Monte Carlo.

2.1.1 Importance Sampling

The principle of importance sampling is to change the probability measure under which samples are drawn. For example, in option pricing, importance sampling aims to transform sample paths that would otherwise end up out-of-the-money into in-the-money ones that will contribute to the derivative price estimation. Importance sampling is particularly useful for rare event simulation.

Let P denote the true derivative price

$$P=\mathbb{E}_f\left[\phi(S)\right]=\int\phi(s)f(s)ds=\int\phi(s)\frac{f(s)}{g(s)}g(s)ds=\mathbb{E}_g\left[\phi(S)\frac{f(S)}{g(S)}\right]$$

where $\phi(s)$ is the payoff and g satisfies $f(s)>0\Rightarrow g(s)>0$. Here, f and g are referred to as the original and the importance sampling densities. The likelihood ratio $\frac{f(s)}{g(s)}$ is the Radon-Nikodym derivative of the original probability measure with respect to the importance sampling probability measure. The importance sampling estimator of the derivative price is given by

$$\hat{P}_{IS}=\frac{1}{n}\sum_{i=1}^{n}\phi(s_i)\frac{f(s_i)}{g(s_i)}$$

where $\phi(s_i)$ is the payoff associated with the i-th sample path. We note that the importance sampling estimator is unbiased and consistent.

The variance of the importance sampling estimator is

$$\begin{aligned}
Var_g\left(\hat{P}_{IS}\right) &= \frac{1}{n}Var\left(\phi(S)\frac{f(S)}{g(S)}\right) \\
&= \frac{1}{n}\left(\int\left(\phi(s)\frac{f(s)}{g(s)}\right)^2 g(s)ds - P^2\right) \\
&= \frac{1}{n}\left(\int\phi(s)^2\frac{f(s)}{g(s)}f(s)ds - P^2\right)
\end{aligned}$$

while the variance of the original estimator is

$$\begin{aligned}
Var_f\left(\hat{P}\right) &= \frac{1}{n}Var\left(\phi(S)\right) \\
&= \frac{1}{n}\left(\int\phi(s)^2 f(s)ds - P^2\right)
\end{aligned}$$

The amount of the variance reduction is given by

$$\frac{1}{n}\left(Var_f\left(\phi(S)\right) - Var_g\left(\phi(S)\frac{f(S)}{g(S)}\right)\right) = \frac{1}{n}\int\phi(s)^2\left(1-\frac{f(s)}{g(s)}\right)f(s)ds$$

If the density g can be chosen so that $\phi(S)\frac{f(S)}{g(S)}$ has a small variance, then importance sampling results in an efficient estimation of P. To make the variance of $\phi(S)\frac{f(S)}{g(S)}$ small, $\frac{f(s)}{g(s)}$ should be inversely related to $\phi(s)f(s)$. If the importance sampling density is not chosen carefully, importance sampling may actually increase the variance. Therefore, to prevent $Var_g\left(\phi(S)\frac{f(S)}{g(S)}\right)$ from exploding it is crucial that $\frac{f(s)}{g(s)}$ is inversely related to $\phi(s)f(s)$.

Suppose that

$$g^*(s) = f(s)\frac{\phi(s)}{P}$$

With such choice of g, $Var_g\left(\phi(S)\frac{f(S)}{g^*(S)}\right) = 0$. Of course this is not feasible in practice as the true derivative price, P, is not known. After all, it is P that we want to estimate. It is suffice to note here that the efficiency of the importance sampling critically depends on the choice of the importance sampling density. Further details of importance sampling are discussed in Glasserman (2003).

2.1.2 Control Variate

The control variate prices one derivative based on the price of another closely related derivative for which an analytical solution exists. Let $\tilde{P}_{CV}$ denote the true price of derivative A that is of interest

$$\tilde{P}_{CV} = \hat{P}_{CV} - \beta\left(\hat{P} - P\right)$$

where $\hat{P}_{CV}$ denotes the Monte Carlo estimate of derivative A, $\hat{P}$ denotes the Monte Carlo estimate of a second derivative B that is used as a control variate and P denotes the price of derivative B obtained using the analytical solution. Derivative B is the control variate, because the observed simulation error $\hat{P} - P$ is used as a control in estimating the price of derivative A. A control variate can be the same derivative priced under different model, different derivative priced under the same model, or different derivative priced under different model. A well known example is the use of a geometric average Asian option as a control variate for pricing arithmetic average Asian option under the Black-Scholes model. We also note that the control variate estimator is unbiased and consistent.

Let σ_{CV}^2 be the variance of $\hat{P}_{CV}$ and σ^2 be the variance of $\hat{P}$. The variance of the control variate estimator is

$$\begin{aligned}
Var\left[\tilde{P}_{CV}\right] &= Var\left[\hat{P}_{CV} - \beta\left(\hat{P} - P\right)\right] \\
&= Var\left[\hat{P}_{CV} - \beta\hat{P} + \beta P\right] \\
&= Var\left[\hat{P}_{CV} - \beta\hat{P}\right] + Var\left[\beta P\right] - 2Cov\left[\hat{P}_{CV} - \beta\hat{P}, \beta P\right] \\
&= \sigma_{CV}^2 + \beta^2\sigma^2 - 2\beta Cov\left[\hat{P}_{CV}, \hat{P}\right] \qquad (2.2)
\end{aligned}$$

The optimal value β^* for which $Var\left[\tilde{P}_{CV}\right]$ attains its minimum is given by

$$\beta^* = \frac{Cov\left[\hat{P}_{CV}, \hat{P}\right]}{\sigma^2}$$

where σ^2 and $Cov\left[\hat{P}_{CV}, \hat{P}\right]$ are estimated using sample variance and sample covariance respectively. β^* is positive (negative) when $\hat{P}_{CV}$ and $\hat{P}$ are positively (negatively) correlated. If $\hat{P}_{CV}$ and $\hat{P}$ are positively correlated, then $\hat{P}_{CV}$ is large (small) when $\hat{P}$ is large (small). Hence, if the realization of $\hat{P}$ is larger (smaller) than P, then it is likely that the realization of $\hat{P}_{CV}$ is also larger (smaller) than P_{CV}. Control variate corrects for this by lowering (raising) $\hat{P}_{CV}$ by $\beta\left(\hat{P} - P\right)$. A similar argument applies when $\hat{P}_{CV}$ and $\hat{P}$ are negatively correlated.

Substituting β^* into (2.2) yields

$$Var\left[\tilde{P}_{CV}\right] = \sigma_{CV}^2 - \frac{Cov\left[\hat{P}_{CV}, \hat{P}\right]^2}{\sigma^2}$$

The effectiveness of the control variate is determined by the correlation between $\hat{P}_{CV}$ and $\hat{P}$. If this correlation is ± 1, then the control variate estimator has zero variance. More information on control variate can be found in Ross (2013).

2.1.3 Antithetic Variates

The main idea behind antithetic variates is to reduce the variance by introducing negatively correlated samples. If ϵ is drawn from the Gaussian distribution, then $-\epsilon$ is also Gaussian; ϵ and $-\epsilon$ are called antithetic variates. Let p_i and p_i' be samples of a payoff obtained from ϵ_i and $-\epsilon_i$ respectively, where $i = 1, \cdots, \frac{n}{2}$. The antithetic estimator of the derivative price is given by

$$\hat{P}_A = \frac{1}{2} \sum_{i=1}^{\frac{n}{2}} \frac{p_i + p_i'}{\frac{n}{2}}$$

Since, for each i, p_i and p_i' have the same distribution, both $\bar{p} = \sum_{i=1}^{\frac{n}{2}} \frac{p_i}{\frac{n}{2}}$ and $\bar{p}' = \sum_{i=1}^{\frac{n}{2}} \frac{p_i'}{\frac{n}{2}}$ are unbiased estimators of the derivative price which implies that the antithetic estimator is also unbiased. We also note that the consistency of the antithetic estimator follows from Slutsky's theorem (Amemiya (1985)).

The variance of the antithetic estimator is

$$Var\left[\hat{P}_A\right] = Var\left[\frac{1}{2}\left(\bar{p} + \bar{p}'\right)\right] = \frac{1}{4}\left(Var\left[\bar{p}\right] + Var\left[\bar{p}'\right] + 2Cov\left[\bar{p}, \bar{p}'\right]\right) \qquad (2.3)$$

Since p_i is obtained from ϵ_i and p_i' is obtained from $-\epsilon_i$, p_i' is likely to be negatively correlated with p_i, resulting in a smaller variance in (2.3). If the payoff function is linear, the variance of the antithetic variates estimator is 0. On the other hand, if the payoff function is symmetric, antithetic variates estimator does not reduce the variance. In general, the more linear the payoff function, the more efficient the antithetic variable estimator. The performance of the antithetic variates estimator depends also on how much extra computational effort is required to obtain antithetic sample compared to thetic sample. For example, if generating antithetic sample requires as much computational effort as generating the thetic sample, then there would be an improvement only if $Corr\left[\bar{p}, \bar{p}'\right] < 0$. On the other hand, if generating antithetic sample requires no extra computational effort compared to generating thetic sample, then there would be an improvement as long as $Corr[\bar{p}, \bar{p}'] < 1$. For a more detailed discussion of antithetic variates see Webber (2011).

2.1.4 Conditional Monte Carlo

The idea behind the conditional Monte Carlo is to compute the derivative price conditioned on the value of the auxiliary random variable.

Suppose that the conditional expectation $\mathbb{E}[\phi(S) \mid Z]$ can be computed analytically, but the distribution of Z is unknown. Let P be the true derivative price. By the tower property of the conditional expectation

$$P = \mathbb{E}[\phi(S)] = \mathbb{E}[\mathbb{E}[\phi(S) \mid Z]]$$

The conditional Monte Carlo estimator is given by

$$\hat{P}_C = \frac{1}{n}\sum_{i=1}^{n} \mathbb{E}[\phi(S) \mid Z_i]$$

We note that the Conditional Monte Carlo estimator is unbiased and consistent.

Recall the law of total variance

$$Var(\phi(S)) = \mathbb{E}[Var[\phi(S) \mid Z]] + Var(\mathbb{E}[\phi(S) \mid Z]) \tag{2.4}$$

Since the conditional Monte Carlo estimator eliminates the variance associated with $Var(\mathbb{E}[\phi(S) \mid Z])$, by (2.4) it has a lower variance compared to the Monte Carlo estimator. Variance reduction can be substantial when S strongly depends on Z. Further discussion of conditional Monte Carlo is in Hirsa (2012).

2.1.5 Stratified Sampling

Stratified sampling is based on dividing the sampling domain into non-overlapping subintervals known as strata and then sampling proportionally from the different strata. Let P denote the true derivative price

$$\begin{aligned} P &= \mathbb{E}\left[\phi\left(S\right)\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\phi\left(S\right) \mid I\right]\right] \\ &= \sum_{i=1}^{k} p_i \mathbb{E}\left[\phi\left(S\right) \mid I = i\right] \\ &= p_1 P_1 + \cdots + p_k P_k \end{aligned}$$

where p_i is the probability that S belongs to the i-th strata, n_i is the number of sample paths in the i-th strata and $n = \sum_{i=1}^{k} n_i$ is the total number of sample paths. I is the stratification variable which partitions the sampling domain into k subintervals according to its value. Here, in contrast to the conditional Monte Carlo, the distribution of the stratification variable is known, but the conditional expectations have to be estimated using Monte Carlo simulation.

The stratified sampling estimator is given by

$$\hat{P}_{SS} = p_1 \hat{P}_1 + \cdots + p_k \hat{P}_k$$

where $\hat{P}_i = \frac{1}{n_i} \sum_{j=1}^{n_i} \phi\left(s_j^i\right)$ and $\phi(s_j^i)$ is the payoff associated with the j-th sample path. We note that the stratified sampling estimator is unbiased and consistent.

The variance of the stratified sampling estimator is

$$\begin{aligned} Var\left(\hat{P}_{SS}\right) &= Var\left(p_1 \hat{P}_1 + \cdots + p_k \hat{P}_k\right) \\ &= \sum_{i=1}^{k} p_i^2 \frac{\sigma_i^2}{n_i} \\ &= \frac{1}{n} \sum_{i=1}^{k} p_i \sigma_i^2 \end{aligned} \tag{2.5}$$

where $p_i = \frac{n_i}{n}$ and $\frac{\sigma_i^2}{n_i} = Var\left(\hat{P}_i\right)$. This approach is called proportional allocation. By the law of total variance, we have $\frac{\sigma^2}{n} \geq \frac{\sum_{i=1}^{k} p_i \sigma_i^2}{n}$ where $\frac{\sigma^2}{n}$ is the variance of an estimator without stratified sampling. However, the optimal approach is to choose n_i in order to minimize the variance

$$\min_{n_i} \sum_{i=1}^{k} p_i^2 \frac{\sigma_i^2}{n_i}$$

$$\text{subject to } \sum_{i=1}^{k} n_i = n.$$

Applying a Lagrange multiplier, the optimal solution is

$$n_i^* = \frac{p_i \sigma_i}{\sum_{i=1}^{k} p_i \sigma_i} n$$

Substituting n_i^* into (2.5) yields

$$Var\left(\hat{P}\right) = \frac{1}{n} \left(\sum_{i=1}^{m} p_i \sigma_i \right)^2$$

Variance reduction due to stratified sampling can be substantial when the value of I strongly affects $\mathbb{E}\left[\phi\left(S\right) \mid I\right]$. For a more detailed overview of stratified sampling see Lemieux (2009).

2.1.6 Multi-Level Monte Carlo

The Multi-Level Monte Carlo was introduced in Giles (2008). It is a Monte Carlo simulation with different number of time steps of size $h_l = 2^{-l}T$ on each level $l = 0, 1, \cdots, L$. For example, when $l = 0$, there is only one time step of size $h_0 = T$. When $l = 1$, there are two times steps each of size $h_1 = \frac{T}{2}$. Finally, when $l = L$, there are 2^L times steps each of size $h_L = \frac{T}{2^L}$.

Let P denote the true derivative price

$$P = \mathbb{E}\left[\phi_L\left(S\right)\right] = \mathbb{E}\left[\phi_0\left(S\right)\right] + \sum_{l=1}^{L} \mathbb{E}\left[\phi_l\left(S\right) - \phi_{l-1}\left(S\right)\right] \tag{2.6}$$

where $\phi_l\left(s\right)$ denotes the payoff on level l. The Multi-Level Monte Carlo estimator is given by

$$\hat{P} = \sum_{l=0}^{L} \hat{P}_l$$

$\hat{P}_0$ is an estimator for $\mathbb{E}\left[\phi_0\left(S\right)\right]$ calculated as a mean of N_l independent sample paths

$$\hat{P}_0 = \frac{1}{N_0} \sum_{i=1}^{N_0} \phi_0\left(s_i\right)$$

and $\hat{P}_l$ is an estimator for $\mathbb{E}\left[P_l - P_{l-1}\right]$ calculated as a mean of N_l independent sample paths

$$\hat{P}_l = \frac{1}{N_l}\sum_{i=1}^{N_l}\left(\phi_l\left(s_i\right) - \phi_{l-1}\left(s_i\right)\right)$$

The variance of the combined Multi-Level Monte Carlo estimator on level l is

$$\begin{aligned} Var\left[\hat{P}_l\right] &= Var\left[\frac{1}{N_l}\sum_{i=1}^{N_l}\left(\phi_l\left(s_i\right) - \phi_{l-1}\left(s_i\right)\right)\right] \\ &= \frac{1}{N_l^2}\sum_{i=1}^{N_l} Var\left[\phi_l\left(s_i\right) - \phi_{l-1}\left(s_i\right)\right] \\ &= \frac{V_l}{N_l} \end{aligned}$$

Thus, the variance of the combined Multi-Level Monte Carlo estimator is

$$Var\left[\hat{P}\right] = \sum_{l=0}^{L} Var\left[\hat{P}_l\right] = \sum_{l=0}^{L}\frac{1}{N_l}V_l$$

where $V_l = Var\left[\phi_l\left(S\right) - \phi_{l-1}\left(S\right)\right]$. Furthermore,

$$\begin{aligned} V_l &= Var\left[\phi_l\left(S\right) - \phi_{l-1}\left(S\right)\right] \\ &= Var\left[\phi_l\left(S\right)\right] + Var\left[\phi_{l-1}\left(S\right)\right] - 2Cov\left[\phi_l\left(S\right), \phi_{l-1}\left(S\right)\right] \end{aligned}$$

so the higher the correlation between $\phi_l\left(S\right)$ and $\phi_{l-1}\left(S\right)$, the lower the variance of the Multi-Level Monte Carlo estimator.

In order to minimize the variance of the Multi-Level Monte Carlo estimator for a given computational cost $C = \sum_{l=0}^{L} N_l \frac{T}{h_l}$, it is possible to use the Lagrange multiplier method. The Lagrangian is given by

$$\mathcal{L} = \sum_{l=0}^{L}\frac{1}{N_l}V_l + \lambda\left(\sum_{l=0}^{L} N_l\frac{T}{h_l} - C\right) \tag{2.7}$$

Differentiating (2.7) with respect to N_l and applying the first order condition shows that the variance is minimized at

$$N_l^* = \sqrt{\frac{V_l h_l}{\lambda T}} \tag{2.8}$$

With such choice of N_l the variance of the combined Multi-Level Monte Carlo estimator becomes

$$Var\left[\hat{P}\right] = \sum_{l=0}^{L} \frac{1}{N_l^*} V_l = \sum_{l=0}^{L} \sqrt{\frac{V_l \lambda T}{h_l}}$$

By Theorem 3.1 in Giles (2008), $Var\left[\hat{P}\right] \leq \frac{\epsilon^2}{2}$ where ϵ is a user-specified accuracy. It follows that

$$\sum_{l=0}^{L} \sqrt{\frac{V_l \lambda T}{h_l}} \leq \frac{\epsilon^2}{2}$$

Substituting $\frac{\sqrt{V_l h_l}}{N_l^*}$ for $\sqrt{\lambda T}$ from (2.8) yields

$$N_l^* \geq 2\epsilon^{-2} \sqrt{V_l h_l} \left(\sum_{l=0}^{L} \sqrt{\frac{V_l}{h_l}} \right)$$

Therefore, the optimal number of sample paths for level l, in order to minimize the variance of the Multi-Level Monte Carlo estimator for a given computational cost C, is

$$N_l^* = \left\lceil 2\epsilon^{-2} \sqrt{V_l h_l} \left(\sum_{l=0}^{L} \sqrt{\frac{V_l}{h_l}} \right) \right\rceil \tag{2.9}$$

Overall, Monte Carlo has computational cost proportional to ϵ^{-3}, whereas that of the Multi-Level Monte Carlo is proportional to $\epsilon^{-2} (\log \epsilon)^2$ due to reduced variance.

Chapter 3

Derivative Pricing and Greeks with Stochastic Volatility Using Importance Sampling

3.1 Introduction

This chapter is motivated by the lack of efficient application of importance sampling to derivative pricing and risk management with price dynamics that have stochastic volatility and jumps. Our contributions are as follows. First, we apply the importance sampling scheme introduced in Fouque and Tullie (2002) to price derivatives under the Heston and the Bates models. To accommodate the jumps in the Bates model, we add a modification so that the effect of jump is captured by the change in volatility of the stock price. Consequently, it is possible to apply importance sampling to the Bates model. Next, we introduce an effective numerical scheme based on a 2-dimensinal lookup table that dramatically improves the speed of this importance sampling scheme. Finally, we demonstrate how the Likelihood Ratio Method for calculating the Greeks can be made more efficient by importance sampling.

It is well known that the error rate in Monte Carlo estimator decreases with $\sqrt{n}$ where n is the number of sample paths. This has an important implication on the speed of the simulation. For example, to reduce the error by a factor of 10, Monte Carlo simulation requires 100 times more samples and thus 100 times more computational effort. Many techniques have been suggested to improve the efficiency of Monte Carlo simulation including: antithetic variates, control variate, importance sampling and stratified sampling. Here, we focus on the variance reduction through importance sampling. The main idea behind importance sampling is to focus the simulation effort on the event of major concern. In option pricing, this amounts to simulating only sample paths that will lead to a non-zero payoff. This is done by changing the drift of the simulated sample paths and then rescaling the payoff by the likelihood ratio to obtain its actual risk-neutral probability. From the mathematical point of view, this likelihood ratio is the Radon-Nikodym derivative of the original probability measure with respect to the importance sampling measure. It is important to note here that the variance reduction achieved through importance sampling critically depends on the change of drift.

The rest of the chapter is organized as follows. In Section 3.2, we discuss how this chapter is set in context with the previous literature. In Section 3.3, we follow Fouque and Tullie (2002) to derive the optimal change of drift for the Heston model. Next, we recall the approximations to European option prices obtained using small noise and fast mean-reversion expansions. Small noise expansion was introduced in Fournie, Lebuchoux and Touzi (1997) and is obtained by regular perturbation of the pricing partial differential equation. Fast mean-reversion expansion was introduced in Fouque and Tullie (2002) and is obtained by singular perturbation of the pricing partial differential equation. We then use these expansions together with the optimal change of drift and apply them to European and barrier options. This results in two approximations to the optimal change of drift. The first approximation to the optimal change of drift is based on the small noise expansion, whereas the second one is based on the fast mean-reversion expansion. Finally, we discuss how to improve the speed of both small noise and fast mean-reversion importance sampling schemes by using a 2-dimensinal lookup table. In Section 3.4, we extend the results in Section 3.3 and derive the optimal change of drift for the Bates model. In Section 3.5, we show how the Greeks can be computed using the Likelihood Ratio Method. We then improve the computational efficiency of estimating the Greeks by using these importance sampling schemes and demonstrate their performance using European and barrier options in both Heston and Bates models. Finally, Section 3.6 concludes the chapter. Throughout this chapter we use a European call and a down-and-out put as examples.

3.2 Literature Review

Importance sampling is a popular variance reduction technique. Much research has been devoted to importance sampling and the choice of drift that maximizes the variance reduction. One of the areas where importance sampling finds fruitful applications is credit risk, because it is able to capture the rare but large losses that are very important in credit risk. For example, Merino and Nyfeler (2004) apply importance sampling in the estimation of individual risk contributions of the obligors to the expected shortfall of the credit portfolio. Glasserman (2005) also uses importance sampling to estimate marginal risk contributions to the portfolio expected shortfall and value-at-risk. Glasserman (2005) notes that individual risk contributions can be thought of as marginal risk contributions and calculated as expected losses conditioned on a tail event. Since calculating these conditional expectations requires simulation of rare events, importance sampling focuses the simulation effort on these rare events thereby dramatically increasing the convergence of the Monte Carlo estimator.

In a copula literature, Glasserman and Li (2005) introduce asymptotically optimal changes of measure for the estimation of credit risk with the Gaussian copula. Glasserman, Kang and Shahabuddin (2008) extend the work of Glasserman and Li (2005) to the multifactor Gaussian copula and Kang and Shahabuddin (2005) to the t-copula. Huang et al. (2010) demonstrate the application of importance sampling to estimate conditional value-at-risk in the Gaussian copula model. Sak and Hörmann (2012) use importance sampling in the same model to estimate the probability of large losses and expected shortfall. Chan and Kroese (2010) utilize importance sampling to estimate the probability of large losses in the t-copula model. Grudke (2009) shows the benefit of using importance sampling for calculating economic capital.

In credit derivatives, Joshi and Kainth (2004) introduce an importance sampling procedure for the pricing of n-th to default swaps. Chen and Glasserman (2008) improve their method by changing the default probabilities and use it for the pricing of basket default swaps. Joshi (2004) shows that importance sampling can be applied to the pricing of single tranche CDOs. In case of default swaps and CDOs the advantage of importance sampling is the concentration of the simulation effort in regions where defaults are most likely.

In option pricing, Vázquez-Abad and Dufresne (1998) introduce a variance reduction method that combines importance sampling and control variate for Asian options in the Black-Scholes model. Glasserman et al. (1999a) propose a variance reduction method for path-dependent options that is based on importance sampling with deterministic change of drift and stratified sampling. They show that the variance of the Monte Carlo estimator can be approximated using the Laplace method for integrals. Based on that, they identify the asymptotically optimal change of drift as a solution to a deterministic optimization problem in discrete time. In a subsequent paper, Glasserman et al. (1999b) apply this method in the Heath-Jarrow-Morton framework. Étoré and Jourdain (2010) improve the method of Glasserman et al. (1999a) by using stratified sampling with adaptive allocation. Guasoni and Robertson (2008) generalize the importance sampling result of Glasserman et al. (1999a) to continuous time. The asymptotically optimal change of drift is defined as a solution to a variational problem that reduces to an Euler-Lagrange equation. Vázquez-Abad and Dufresne (1998), Su and Fu (1999), as well as Arouna (2004) use stochastic approximation methods to determine the optimal change of drift. In the least squares importance sampling of Caprotti (2008), the change of drift is formulated in terms of least squares minimization problem. As shown in Caprotti (2007), least squares importance sampling can be also used for LIBOR market model. The common feature of these studies is that the optimal change of drift is obtained by solving an optimization problem.

Fouque and Tullie (2002) develop a variance reduction method based on importance sampling with a stochastic change of drift for pricing options under stochastic volatility. The change of drift is based on the approximation of the option price obtained by either fast mean-reversion or small noise expansions. Fast mean-reversion expansion is the approximation to the option price given by the Black-Scholes price with long-run average volatility in the leading term. Small noise expansion is another approximation to the option price given by the Black-Scholes price with initial volatility. Fast mean-reversion expansion works well for mean-reversion rate being large and small noise expansion works well for mean-reversion rate being small. Fouque and Tullie (2002) apply importance sampling using fast mean-reversion and small noise expansion for European option in the Scott's stochastic volatility model and Tullie (2002) also to American and barrier options. Importance sampling for pricing barrier options under the Black-Scholes model and LIBOR market model was used by Glasserman and Staum (2001). Joshi and Leung (2007) also apply importance sampling to speed up pricing of barrier options in the Merton jump-diffusion model. Fouque and Han (2004) apply importance sampling to European as well as Asian options in a 2-factor stochastic volatility model. Robertson (2010) considers the Heston model and uses Sample Path Large Deviation Principles to identify an asymptotically optimal change of drift for Asian options that is deterministic. It is defined as a solution to a 2-dimensional variational problem that reduces to an Euler-Lagrange equation. This is very much in the spirit of Guasoni and Robertson

(2008), with the difference that the variational problem is 2-dimensional and the corresponding Euler-Lagrange equation is more complicated.

3.3 Heston Model with Importance Sampling

The dynamics of the Heston model under the risk-neutral measure $\mathbb{Q}$ is given by

$$dS_t = rS_t dt + \sqrt{v_t} S_t dW_t^S \tag{3.1}$$

$$dv_t = \kappa(\theta - v_t)\, dt + \xi \sqrt{v_t} dW_t^v \tag{3.2}$$

where $\left\langle dW_t^S, dW_t^v \right\rangle = \rho dt$. S_t is the stock price, r is the risk-free interest rate, v_t is the variance, κ is the mean-reversion rate, θ is the long-term variance, ξ is the volatility of volatility and ρ is the correlation between stock returns and changes in the variance.

In matrix notation, the dynamics in (3.1) and (3.2) can be represented as

$$dX_t = b(X_t)\, dt + a(X_t)\, d\eta_t \tag{3.3}$$

where η_t is a 2-dimensional $\mathbb{Q}$-Brownian motion and

$$X_t = \begin{pmatrix} S_t \\ v_t \end{pmatrix}$$

$$b(x) = \begin{pmatrix} rs \\ \kappa(\theta - v) \end{pmatrix}$$

$$a(x) = \begin{pmatrix} \sqrt{v}s & 0 \\ \xi\sqrt{v}\rho & \xi\sqrt{v(1-\rho^2)} \end{pmatrix}$$

$$\eta_t = \begin{pmatrix} Z_t^1 \\ Z_t^2 \end{pmatrix}$$

where Z_t^1 and Z_t^2 are independent Brownian motions.

In this setup, the price $P(t, X_t)$ of a European option at time t is given by

$$P(t,x) = \mathbb{E}\left[e^{-r(T-t)} \max(S_T - K) \,|X_t = x\right]$$

Here T is the option maturity, r is the risk-free interest rate, K is the strike price and S_T is the price of the underlying asset at T.

3.3.1 Changing Measure

Following Fouque and Tullie (2002), we derive the optimal change of drift for the Heston model. First, we introduce the martingale

$$H_t = \exp\left(\int_0^t h^\top (s, X_s)\, d\eta_s + \frac{1}{2}\int_0^t \|h(s, X_s)\|^2\, ds\right) \tag{3.4}$$

where $h^\top$ denotes the transpose of h. Next, we define a new probability measure denoted by $\tilde{\mathbb{Q}}$ which is equivalent to $\mathbb{Q}$ by its Radon-Nikodym derivative

$$\frac{d\tilde{\mathbb{Q}}}{d\mathbb{Q}} = (H_T)^{-1}$$

By the Girsanov theorem, the process

$$\tilde{\eta}_t = \eta_t + \int_0^t h(s, X_s)\, ds$$

is a 2-dimensional $\tilde{\mathbb{Q}}$-Brownian motion. Using $\tilde{\eta}_t$, (3.3) and (3.4) can be written as

$$\begin{aligned}
dX_t &= (b(X_t) - a(X_t)\, h(t, X_t))\, dt + a(X_t)\, d\tilde{\eta}_t \\
H_t &= \exp\left(\int_0^t h^\top(s, X_s) d\tilde{\eta}_s - \frac{1}{2}\int_0^t ||h(s, X_s)||^2 ds\right)
\end{aligned}$$

With respect to the new measure, the price of a European option becomes

$$P(t, x) = \tilde{\mathbb{E}}\left[e^{-r(T-t)} \max(S_T - K) H_T | X_t = x\right] \tag{3.5}$$

Next, we apply Ito's lemma to $P(t, X_t)H_t$ and use Kolmogorov backward equation for $P(t, X_t)$

$$\begin{aligned}
d(P(t, X_t)H_t) &= P(t, X_t)dH_t + H_t dP(t, X_t) + H_t h^\top a\nabla P dt \\
&= PH_t h^\top d\tilde{\eta}_t + H_t\left(\frac{\partial P}{\partial t}dt + \nabla P^\top dX + \frac{1}{2}\mathrm{tr}(\nabla^2 P \cdot aa^\top)dt\right) + \\
&+ H_t h^\top a\nabla P dt \\
&= PH_t h^\top d\tilde{\eta}_t + H_t\frac{\partial P}{\partial t}dt + H_t \nabla P^\top\left[(b - ah)\, dt + ad\tilde{\eta}_t\right] + \\
&+ \frac{1}{2}H_t \mathrm{tr}(\nabla^2 P \cdot aa^\top)dt + H_t h^\top a\nabla P dt
\end{aligned}$$

$$
\begin{aligned}
&= PH_t h^\top d\tilde{\eta}_t + H_t \frac{\partial P}{\partial t} dt + H_t \nabla P^\top b dt + H_t \nabla P^\top a d\tilde{\eta}_t + \\
&+ \frac{1}{2} H_t \mathrm{tr}(\nabla^2 P \cdot aa^\top) dt \\
&= PH_t h^\top d\tilde{\eta}_t + H_t \nabla P^\top a d\tilde{\eta}_t \\
&= H_t \left(\nabla P^\top a + Ph^\top \right) d\tilde{\eta}_t
\end{aligned}
$$

where ∇P is the gradient of P with respect to the state variable x. Integrating $d\left(P(t, X_t)H_t\right)$ from 0 to T yields

$$
P(T, X_T)H_T = P(0, X_0)H_0 + \int_0^T H_t \left(P(t, X_t)h(t, X_t)^\top + \nabla P(t, X_t)^\top a(t, X_t) \right) d\tilde{\eta}_t
$$

which reduces to

$$
\phi(X_T)H_T = P(0, x) + \int_0^T H_t \left(P(t, X_t)h(t, X_t)^\top + \nabla P(t, X_t)^\top a(t, X_t) \right) d\tilde{\eta}_t
$$

where $\phi(X_T)$ is the option payoff at maturity. Therefore, the variances of Monte Carlo estimators under $\tilde{\mathbb{Q}}$ and $\mathbb{Q}$, respectively, are given by

$$
\begin{aligned}
Var_{\tilde{\mathbb{Q}}}\left(\phi(X_T)H_T\right) &= \tilde{\mathbb{E}}\left[\int_0^T H_t^2 ||P(t, X_t)h(t, X_t)^\top + \nabla P(t, X_t)^\top a(t, X_t)||^2 dt \right] \\
Var_{\mathbb{Q}}\left(\phi(X_T)\right) &= \mathbb{E}\left[\int_0^T ||\nabla P(t, X_t)^\top a(t, X_t)||^2 dt \right]
\end{aligned}
$$

and the optimal choice of h for which the variance of $\phi(X_T)H_T$ under $\tilde{\mathbb{Q}}$ is minimized is

$$
h(t, X_t) = -\frac{1}{P(t, X_t)} a(t, X_t)^\top \nabla P(t, X_t) \tag{3.6}
$$

The difficulty with (3.6) is that $P(t, X_t)$ and $\nabla P(t, X_t)$ are not known. Indeed, the whole purpose of the simulation is to find $P(t, X_t)$. Hence, the Black-Scholes equivalents of $P(t, X_t)$ and $\nabla P(t, X_t)$ will be used instead. In what follows, we study the conditions under which the Heston model can be approximated by the Black-Scholes model.

Following Fouque, Papanicolaou and Sircar (2011), small noise expansion of the Heston model is obtained by setting $\kappa = \delta$ and $\xi = m\sqrt{2\delta}$. The dynamics of the Heston model becomes

$$
\begin{aligned}
dS_t &= rS_t dt + \sqrt{v_t} S_t dW_t^S \\
dv_t &= \delta\left(\theta - v_t\right) dt + m\sqrt{2\delta v_t} dW_t^v
\end{aligned}
$$

where $\langle dW_t^S, dW_t^v \rangle = \rho dt$, δ is a small positive parameter and the Feller condition, $m^2 < \theta$, ensures that v_t is always positive.

This leads to the following fundamental partial differential equation for pricing derivatives

$$\begin{aligned} 0 &= \frac{\partial P}{\partial t} + \frac{1}{2} v S^2 \frac{\partial^2 P}{\partial S^2} + \rho m \sqrt{2\delta} v S \frac{\partial^2 P}{\partial S \partial v} + m^2 \delta v \frac{\partial^2 P}{\partial v^2} \\ &+ r \left(S \frac{\partial P}{\partial S} - P \right) + \left[\delta \left(\theta - v \right) - \lambda v \right] \frac{\partial P}{\partial v} \end{aligned}$$

which can be written as

$$\left(\delta \mathcal{M}_0 + \sqrt{\delta} \mathcal{M}_1 + \mathcal{M}_2 \right) P = 0 \tag{3.7}$$

with differential operators

$$\begin{aligned} \mathcal{M}_0 &= m^2 v \frac{\partial^2}{\partial v^2} + (\theta - v) \frac{\partial}{\partial v} \\ \mathcal{M}_1 &= \rho m \sqrt{2} v S \frac{\partial^2}{\partial S \partial v} \\ \mathcal{M}_2 &= \frac{\partial}{\partial t} + \frac{1}{2} v S^2 \frac{\partial^2}{\partial S^2} + r \left(S \frac{\partial}{\partial S} - . \right) - \lambda v \frac{\partial}{\partial v} \end{aligned}$$

When the market price of volatility risk λ is 0, $\mathcal{M}_2$ is the Black-Scholes operator with volatility $\sqrt{v}$. If $\delta = 0$, (3.7) becomes

$$\mathcal{M}_2 P = 0$$

Therefore, the approximation P_{SNE} of $P(t, X_t)$ is given by the Black-Scholes formula with volatility $\sqrt{v}$. When volatility is slowly varying, that is when $\kappa \to 0$, volatility is stuck at its initial level. In this case, the approximation P_{SNE} of $P(t, X_t)$ is given by the Black-Scholes formula with volatility $\sqrt{v_0}$.

As an alternative, the fast mean-reversion expansion of the Heston model is obtained by setting $\kappa = \frac{1}{\varepsilon}$ and $\xi = \frac{m\sqrt{2}}{\sqrt{\varepsilon}}$. The dynamics of the Heston model becomes

$$\begin{aligned} dS_t &= r S_t dt + \sqrt{v_t} S_t dW_t^S \\ dv_t &= \frac{1}{\varepsilon} \left(\theta - v_t \right) dt + \frac{m}{\sqrt{\varepsilon}} \sqrt{2 v_t} dW_t^v \end{aligned}$$

where $\langle dW_t^S, dW_t^v \rangle = \rho dt$, ε is a small positive parameter and the Feller condition, $m^2 < \theta$, ensures that v_t is always positive.

Similarly, the fundamental partial differential equation for pricing derivatives is

$$
\begin{aligned}
0 &= \frac{\partial P}{\partial t} + \frac{1}{2} v S^2 \frac{\partial^2 P}{\partial S^2} + \rho \frac{m\sqrt{2}}{\sqrt{\varepsilon}} v S \frac{\partial^2 P}{\partial S \partial v} + \frac{m^2}{\varepsilon} v \frac{\partial^2 P}{\partial v^2} \\
&+ r\left(S\frac{\partial P}{\partial S} - P\right) + \left[\frac{1}{\varepsilon}(\theta - v) - \lambda v\right]\frac{\partial P}{\partial v}
\end{aligned}
$$

which can be written as

$$
\left(\frac{1}{\varepsilon}\mathcal{L}_0 + \frac{1}{\sqrt{\varepsilon}}\mathcal{L}_1 + \mathcal{L}_2\right) P = 0 \tag{3.8}
$$

with differential operators

$$
\begin{aligned}
\mathcal{L}_0 &= m^2 v \frac{\partial^2}{\partial v^2} + (\theta - v)\frac{\partial}{\partial v} \\
\mathcal{L}_1 &= \rho m \sqrt{2} v S \frac{\partial^2}{\partial S \partial v} \\
\mathcal{L}_2 &= \frac{\partial}{\partial t} + \frac{1}{2} v S^2 \frac{\partial^2}{\partial S^2} + r\left(S\frac{\partial}{\partial S} - .\right) - \lambda v \frac{\partial}{\partial v}
\end{aligned}
$$

Expanding P in powers of $\sqrt{\varepsilon}$, we get

$$
P = P_0 + \sqrt{\varepsilon} P_1 + \cdots
$$

Substituting this into (3.8) yields

$$
\begin{aligned}
\frac{1}{\varepsilon}\mathcal{L}_0 P_0 &+ \frac{1}{\sqrt{\varepsilon}}\left(\mathcal{L}_0 P_1 + \mathcal{L}_1 P_0\right) \\
&+ \left(\mathcal{L}_2 P_0 + \mathcal{L}_1 P_1\right) \\
&+ \sqrt{\varepsilon}\mathcal{L}_2 P_1 \\
&+ \cdots \\
&= 0
\end{aligned}
$$

We eliminate terms of order $\frac{1}{\varepsilon}$, $\frac{1}{\sqrt{\varepsilon}}$, $\sqrt{\varepsilon}$ by equating them to 0. This leads to

$$
\left(\mathcal{L}_2 P_0 + \mathcal{L}_1 P_1\right) = 0
$$

Since the operator $\mathcal{L}_1$ takes derivatives with respect to v and P_1 must be constant with respect to v, we must have $\mathcal{L}_1 P_1 = 0$. Finally, we have

$$
\mathcal{L}_2 P_0 = 0
$$

When the market price of volatility risk λ is 0, $\mathcal{L}_2$ is the Black-Scholes operator with constant volatility. As shown in Fouque, Papanicolaou and Sircar (2011), $P_0 = P_{BS(\bar{\sigma})}$ where $\bar{\sigma}$ is the constant volatility $\bar{\sigma} = \sqrt{\theta}$. Therefore, the approximation P_{FMR} of $P(t, X_t)$ is given by the Black-Scholes formula with volatility $\sqrt{\theta}$.

From (3.6), under the small noise expansion, h is given by

$$h = -\frac{1}{P_{SNE}} \begin{pmatrix} s\sqrt{v}\frac{\partial P_{SNE}}{\partial s} \\ 0 \end{pmatrix}$$

where P_{SNE} is the option price under the geometric Brownian motion dynamics with volatility $\sqrt{v_0}$. Similarly, under the fast mean-reversion expansion, h is given by

$$h = -\frac{1}{P_{FMR}} \begin{pmatrix} s\sqrt{v}\frac{\partial P_{FMR}}{\partial s} \\ 0 \end{pmatrix}$$

where P_{FMR} is the option price under the geometric Brownian motion dynamics with volatility $\sqrt{\theta}$.

To simulate the Heston model using importance sampling we use the following discretization

$$\begin{aligned} S_{i+1} &= S_i \exp\left(\left(r - \frac{v_i}{2} - \sqrt{v_i} h\left(i, S_i, v_i\right)\right)\Delta t + \sqrt{v_i \Delta t} Z^1_{i+1}\right) \\ v_{i+1} &= v_i + \left(\kappa\left(\theta - v_i\right) - \sqrt{v_i}\xi\rho h\left(i, S_i, v_i\right)\right)\Delta t \\ &\quad + \xi\sqrt{v_i \Delta t}\left(\rho Z^1_{i+1} + \sqrt{1-\rho^2} Z^2_{i+1}\right) \end{aligned}$$

where $\Delta t = t_{i+1} - t_i$, Z^1_{i+1} and Z^2_{i+1} are standard normal random variables. Then from (3.5) the option payoff is multiplied by the Radon-Nikodym derivative of the risk-neutral probability measure with respect to the importance sampling measure

$$H_T = \exp\left(\sum_{i=1}^{M} h\left(i, S_i, v_i\right)\sqrt{\Delta t} Z^1_{i+1} - \frac{1}{2}\sum_{i=1}^{M} h\left(i, S_i, v_i\right)^2 \Delta t\right)$$

where M is the number of time steps.

3.3.2 Numerical Examples

In this section, we present the numerical results for a European call, where a semi-analytical solution exists and for a down-and-out put, where an analytical solution is not available. Here, we set $\rho = -0.4$, $\xi = 0.4$, $v_0 = 0.04$, $\theta = 0.09$,

$S_0 = 100$, $r = 5\%$ and $T = 1$ year. The barrier for the down-and-out put is set at 50. Initial variance of 0.04 corresponds to volatility of 20%, whereas long-run variance of 0.09 corresponds to volatility of 30%. We tested four rates of mean reversion, κ: 0.5, 2, 5, 10 and five strikes, K: 60, 80, 100, 120, 140. As benchmarks, we compared option prices simulated under the importance sampling (Small Noise Expansion, SNE and Fast Mean-Reversion, FMR) against the basic Monte Carlo (MC), control variate (CV) and semi-analytical solution (H). All simulations are performed using the same sequence of pseudo-random numbers. We simulate 10,000 sample paths. For the European call we use a time-increment of 0.001, whereas for the down-and-out put we use a time-increment of 1/252, which corresponds to one business day. As the control variate, we use the geometric Brownian motion dynamics with a volatility of 30% which is the long-run mean of volatility used for the Heston model.

Table 3.1 presents the price, variance and relative error of the European call. The relative error is measured against the semi-analytical solution. Table 3.2 presents the price and variance for the down-and-out put option, where analytical solution is not available. For the European call, the importance sampling strongly outperforms the basic MC in terms of variance reduction and price accuracy. Variance reduction for SNE is up to 105 times compared to MC. For FMR, this is up to 880 times, whereas for CV this is up to 35 times. The amount of variance reduction depends on mean-reversion rate and strike. Greater variance reduction is achieved with higher mean-reversion rate and the more the option is in-the-money. In terms of relative error, SNE tends to exhibit the highest relative error, which may have to do with the parameters of the Heston model or the random number seed. For the European call, CV has comparable performance as the importance sampling; it is more efficient than SNE but inferior to FMR. However, for the barrier option, the superiority of the importance sampling continues, but CV starts to deteriorate especially for in-the-money put option. Variance reduction for SNE is up to 12 times compared to MC. For FMR, this is up to 65 times, whereas for CV this is up to 10 times. As the mean-reversion rate increases and as the option becomes deeper in-the-money, so does the amount of variance reduction. Since there is no analytical solution for the price of a barrier option under the Heston model, it is not possible to compute the relative error and conclude on price accuracy. Between the two importance sampling schemes, FMR is more efficient than SNE even for cases when mean-reversion speed is slow. Fouque and Tullie (2002) report a similar finding for the Scott's model. FMR appears also to be more accurate than SNE.

κ	K	Price					Variance				Relative error (%)			
		MC	SNE	FMR	CV	H	MC	SNE	FMR	CV	MC	SNE	FMR	CV
0.5	60	43.07	43.09	43.16	43.06	43.19	443.13	2.31	2.58	122.19	0.28	0.24	0.08	0.30
0.5	80	25.40	25.23	25.41	25.36	25.46	356.17	18.07	10.87	107.15	0.23	0.89	0.20	0.39
0.5	100	10.92	10.63	10.88	10.88	10.92	206.48	33.42	20.99	77.83	0.05	2.66	0.41	0.33
0.5	120	2.99	2.79	2.99	3.00	3.02	76.87	20.45	11.68	45.09	0.79	7.37	0.89	0.66
0.5	140	0.75	0.64	0.73	0.75	0.75	24.19	10.00	3.10	19.33	0.34	15.19	2.08	0.66
2	60	43.15	43.16	43.20	43.14	43.23	626.72	7.09	1.49	73.43	0.18	0.16	0.07	0.20
2	80	25.94	25.74	25.92	25.89	25.96	509.03	24.22	7.59	65.11	0.06	0.87	0.15	0.27
2	100	12.59	12.27	12.55	12.55	12.58	306.51	42.08	13.27	49.47	0.10	2.44	0.24	0.25
2	120	4.73	4.51	4.77	4.74	4.78	130.56	25.79	8.90	31.96	0.95	5.57	0.16	0.81
2	140	1.46	1.35	1.50	1.48	1.50	43.65	11.78	3.03	17.78	2.24	9.91	0.09	1.07
5	60	43.18	43.18	43.22	43.18	43.23	766.17	4.48	0.87	41.50	0.11	0.12	0.04	0.13
5	80	26.29	26.13	26.27	26.23	26.28	627.81	35.57	4.24	36.92	0.03	0.57	0.06	0.21
5	100	13.56	13.27	13.52	13.51	13.52	393.03	55.45	6.62	28.73	0.26	1.91	0.05	0.13
5	120	5.86	5.66	5.91	5.87	5.91	188.13	38.31	4.59	18.82	0.72	4.20	0.04	0.58
5	140	2.21	2.10	2.26	2.23	2.25	73.74	17.58	1.84	11.16	1.91	6.87	0.26	0.71
10	60	43.20	43.19	43.21	43.19	43.22	837.57	4.00	0.48	23.74	0.05	0.08	0.02	0.07
10	80	26.42	26.28	26.38	26.36	26.39	691.21	36.04	2.05	21.28	0.13	0.40	0.02	0.12
10	100	13.95	13.69	13.89	13.89	13.89	442.58	62.10	2.87	16.77	0.41	1.44	0.01	0.00
10	120	6.36	6.19	6.40	6.37	6.40	224.75	48.08	1.90	11.05	0.53	3.28	0.05	0.38
10	140	2.61	2.51	2.65	2.64	2.65	96.87	26.70	0.77	6.52	1.49	5.17	0.16	0.27

Table 3.1: Basic Monte Carlo (MC), Small Noise Expansion (SNE), Fast Mean-Reversion (FMR) and Control Variate (CV) price, variance of price and relative error for a European call under the Heston model with different mean-reversion rates (κ) and strikes (K). Relative error is measured against the semi-analytical solution (H).

κ	K	Price				Variance			
		MC	SNE	FMR	CV	MC	SNE	FMR	CV
0.5	60	0.06	0.02	0.03	0.06	0.31	0.03	0.11	0.30
0.5	80	1.04	0.81	0.95	1.02	14.67	5.95	4.29	10.47
0.5	100	5.15	4.87	4.99	5.06	91.72	33.69	13.09	44.49
0.5	120	15.89	15.56	15.69	15.75	241.87	44.03	40.04	91.40
0.5	140	32.26	31.98	32.02	32.09	359.81	51.44	39.98	148.68
2	60	0.08	0.03	0.05	0.08	0.42	0.21	0.03	0.37
2	80	1.47	1.25	1.37	1.43	19.73	15.21	1.76	9.83
2	100	6.70	6.42	6.55	6.58	116.03	42.20	9.97	34.99
2	120	17.56	17.25	17.28	17.39	296.60	60.26	22.62	71.92
2	140	32.89	32.57	32.58	32.68	472.43	61.23	32.29	119.85
5	60	0.08	0.05	0.07	0.08	0.43	0.36	0.01	0.34
5	80	1.80	1.63	1.71	1.76	23.35	26.05	0.98	7.94
5	100	7.70	7.45	7.51	7.57	131.74	70.67	5.29	26.91
5	120	18.76	18.45	18.47	18.57	327.30	74.11	12.60	56.02
5	140	33.76	33.37	33.41	33.53	529.30	66.55	21.71	93.07
10	60	0.09	0.06	0.08	0.09	0.43	0.89	0.01	0.26
10	80	1.95	1.80	1.87	1.90	24.96	27.15	0.58	5.07
10	100	8.17	7.91	7.99	8.02	138.33	74.28	3.26	17.00
10	120	19.37	19.05	19.11	19.17	339.36	84.77	8.40	35.38
10	140	34.32	33.96	33.96	34.07	549.13	72.73	15.74	58.47

Table 3.2: Basic Monte Carlo (MC), Small Noise Expansion (SNE), Fast Mean-Reversion (FMR) and Control Variate (CV) price and variance of price for a down-and-out put under the Heston model with different mean-reversion rates (κ) and strikes (K).

So far, we show that importance sampling outperformed both MC and CV in terms of variance reduction especially for barrier option. However, SNE and FMR are computationally expensive, as the change of the drift depends on option price and delta in the Black-Scholes model at every time step. For example, if each sample path is simulated using 1,000 time steps and there are 10,000 sample paths, one has to calculate 10,000,000 Black-Scholes option prices and deltas. There is no such calculation in MC, whereas CV requires the simulation of 10,000 sample paths to obtain the Monte Carlo estimate of the option price used as a control variate. Table 3.3 reports price, variance of price and computational time for an at-the-money European call option, when the mean-reversion is 0.5. Variance reduction ratio and speed are measured against the basic Monte Carlo.

	Price	Variance	Time (sec)	Variance reduction ratio	Speed	Effective performance
MC	10.92	206.48	0.51	1	1	1
CV	10.88	77.83	0.71	2.65	1.39	1.91
SNE	10.63	33.42	62.62	6.18	123.05	0.05
FMR	10.88	20.99	62.27	9.84	122.36	0.08
SNE with lookup table	10.58	37.96	1.18	5.44	2.32	2.34
FMR with lookup table	10.86	21.8	1.23	9.47	2.41	3.93

Table 3.3: Performance measures for basic Monte Carlo (MC), Control Variate (CV), Small Noise Expansion (SNE) and Fast Mean-Reversion (FMR).

Effective performance is defined as the variance reduction ratio over speed. Speed itself is defined as the ratio of computational time of the importance sampling to computational time of the basic Monte Carlo.

We note from Table 3.3 that FMR reduces the variance almost 10 times compared to MC, but it is 120 times slower. SNE reduces the variance 6 times compared to MC, but it is also 120 times slower. Therefore, the effective performance of both FMR and SNE is far worse than that of MC. On the other hand, the effective performance of CV is almost 2 times better than MC.

To improve the effective performance of FMR and SNE, we propose the use of a lookup table. As mentioned before, when calculating the drift in the simulation we repeatedly re-calculate the Black-Scholes option price and delta at every time step for different pairs of stock price and time to maturity. Therefore, it is possible to construct a 2-dimensional lookup table using the range of possible values of the stock price and the time to maturity according to the time step used in the simulation. This way the number of times Black-Scholes option prices and deltas are calculated is significantly reduced.

To demonstrate the performance of this numerical scheme, we used a 2-dimensional lookup table with 451 stock prices [50, 51, ..., 500] and 1000 times to maturity [1, 0.999, ..., 0.001], which corresponds to the 1,000 time steps in the Monte Carlo simulation. Table 3.3 shows that using the lookup table slightly increased the variance of SNE and FMR. As reported in Table 3.3, the use of a lookup table speeds up SNE and FMR by 60 times. Note that the computational time reported in Table 3.3 for SNE and FMR with a lookup table includes time needed to generate the lookup table. It is clear that the lookup table improves the effective performance of SNE and FMR to the point that not only they significantly outperform MC but also CV.

3.4 Adding Jumps

Here, we demonstrate how the importance sampling can still be used when a jump process is added to the dynamics in (3.1). To this end, we follow the stochastic volatility with jump dynamics in Bates (1996). The dynamics of the Bates model under the risk-neutral measure denoted by $\mathbb{Q}$ is given by

$$\begin{aligned} dS_t &= S_t \left(r - \lambda \bar{k}\right) dt + S_t \sqrt{v_t} dW_t^S + S_t dZ_t && (3.9) \\ dv_t &= \kappa \left(\theta - v_t\right) dt + \xi \sqrt{v_t} dW_t^v && (3.10) \end{aligned}$$

where $\left\langle dW_t^S dW_t^v \right\rangle = \rho dt$. S_t is the stock price, r is the risk-free interest rate, v_t is the variance, κ is the mean-reversion rate, θ is the long-term variance, ξ is the volatility of volatility and ρ is the correlation between stock returns and changes in the variance. Z_t is a compound Poisson process with intensity λ and log-normal distribution of jump sizes such that if k is its jump size then $\ln(1+k) \sim \mathcal{N}\left(\ln\left(1+\bar{k}\right) - \frac{\delta^2}{2}, \delta^2\right)$.

For the purpose of importance sampling, we will rewrite the model as follows

$$\begin{aligned} dS_t &= S_t \left(r - \lambda \bar{k}\right) dt + S_t \left(\sqrt{v_t} + \frac{dZ_t}{dW_t^S}\right) dW_t^S && (3.11) \\ dv_t &= \kappa \left(\theta - v_t\right) dt + \xi \sqrt{v_t} dW_t^v && (3.12) \end{aligned}$$

Comparing (3.9) and (3.11), it is clear that the dynamics is the same.

In matrix notation the model dynamics is

$$dX_t = b\left(X_t\right) dt + a\left(X_t\right) d\eta_t \qquad (3.13)$$

where η_t is a 2-dimensional $\mathbb{Q}$-Brownian motion and

$$\begin{aligned} X_t &= \begin{pmatrix} S_t \\ v_t \end{pmatrix} \\ b(x) &= \begin{pmatrix} s\left(r - \lambda \bar{k}\right) \\ \kappa(\theta - v) \end{pmatrix} \\ a(x) &= \begin{pmatrix} s\left(\sqrt{v} + \frac{dZ_t}{dW_t^S}\right) & 0 \\ \xi\sqrt{v}\rho & \xi\sqrt{v\left(1-\rho^2\right)} \end{pmatrix} \end{aligned}$$

$$\eta_t = \begin{pmatrix} Z_t^1 \\ Z_t^2 \end{pmatrix}$$

where Z_t^1 and Z_t^2 are independent Brownian motions.

3.4.1 Importance Sampling

Using the analogous derivation to that presented in Section 3.3, the optimal choice of h for which the variance is minimized has the same form as (3.6). With the Bates model in (3.13) and from (3.6), under the small noise expansion, h is given by

$$h = -\frac{1}{P_{SNE}} \begin{pmatrix} s\left(\sqrt{v} + \frac{dZ}{dW^S}\right) \frac{\partial P_{SNE}}{\partial s} \\ 0 \end{pmatrix} \tag{3.14}$$

where P_{SNE} is the option price under the geometric Brownian motion dynamics with volatility $\sqrt{v_0}$. Similarly, under the fast mean-reversion expansion, h is given by

$$h = -\frac{1}{P_{FMR}} \begin{pmatrix} s\left(\sqrt{v} + \frac{dZ}{dW^S}\right) \frac{\partial P_{FMR}}{\partial s} \\ 0 \end{pmatrix} \tag{3.15}$$

where P_{FMR} is the option price under the geometric Brownian motion dynamics with volatility $\sqrt{\theta}$.

To simulate the Bates model using importance sampling we use the following discretization

$$\begin{aligned} S_{i+1} &= S_i \exp\left(\mu\Delta t + \sqrt{v_i \Delta t} Z_{i+1}^1 + \sum_{j=1}^{N_i} k_j\right) \\ \mu &= \left(r - \lambda\bar{k} - \frac{v_i}{2} - \sqrt{v_i}\left(h(i, S_i, v_i) + \frac{\sum_{j=1}^{N_i} k_j}{\sqrt{\Delta t} Z_{i+1}^1}\right)\right) \\ v_{i+1} &= v_i + \left(\kappa(\theta - v_i) - \sqrt{v_i}\xi\rho h(i, S_i, v_i)\right)\Delta t \\ &\quad + \xi\sqrt{v_i \Delta t}\left(\rho Z_{i+1}^1 + \sqrt{1-\rho^2} Z_{i+1}^2\right) \end{aligned}$$

$\Delta t = t_{i+1} - t_i$, Z_{i+1}^1 and Z_{i+1}^2 are standard normal random variables, N_i is a Poisson random variable with mean $\lambda\Delta t$ and $\ln(1+k) \sim \mathcal{N}\left(\ln\left(1+\bar{k}\right) - \frac{\delta^2}{2}, \delta^2\right)$. If W_{i+1}^S happens to be 0, we set $\frac{\sum_{j=1}^{N_i} k_j}{\sqrt{\Delta t} W_{i+1}^S}$ to 0, which means that we apply importance sampling only to the stochastic volatility component. Then from (3.5) the option payoff is multiplied by the Radon-Nikodym derivative of the risk-neutral probability measure with respect to the importance sampling measure. Here H_T has the same form as (3.6), but h follows that in (3.14) for SNE scheme and (3.15) for FMR scheme.

3.4.2 Numerical Examples

As in the previous section, the price is obtained using five methods: basic Monte Carlo (MC), Small Noise Expansion (SNE), Fast Mean-Reversion (FMR), Control Variate (CV) and semi-analytical solution (B). We assume that the jump intensity is 1 jump per year, standard deviation of the jumps is 2% and the mean jump size is -5%. All other parameters are the same as in the Heston model in Section 3.3.2. As the control variate, we use the geometric Brownian motion dynamics with a volatility of 30%.

Table 3.4 presents the price, variance and relative error for the European call. Relative error is measured against the semi-analytical solution. Importance sampling outperforms CV in terms of both variance reduction and relative error. Variance reduction for SNE is up to 27 times compared to MC. For FMR, this is up to 60 times, whereas for CV this is up to 17 times. As before, the amount of variance reduction depends on mean-reversion rate and moneyness. As the mean-reversion rate increases and as the option becomes deeper in-the-money, so does the amount of variance reduction. FMR appears to be not only more efficient but also more accurate than SNE as measured by the relative error.

Table 3.5 reports the price and variance for the down-and-out put option. The price is obtained using four methods: basic Monte Carlo (MC), Small Noise Expansion (SNE), Fast Mean-Reversion (FMR) and Control Variate (CV). The barrier is set at 50 and we use time increment of 1/252, which corresponds to one business day. All the other parameters are the same as those for the European call. There is no analytical solution for a barrier option price under the Bates model. Since we do not have the analytic solution, we are not able to conclude on price accuracy. However, it is clear from Table 3.5 that importance sampling outperforms CV in terms of variance reduction. Variance reduction for SNE is up to 11 times compared to MC. For FMR, this is up to 42 times, whereas for CV this is up to 6 times. As the mean-reversion rate increases and as the option becomes deeper in-the-money, so does the amount of variance reduction. For the barrier option, we note that the performance of SNE starts to be less competitive especially for fast mean-reversion and deep in-the-money cases, whereas FMR continues to dominate CV.

κ	K	Price					Variance				Relative error (%)			
		MC	SNE	FMR	CV	B	MC	SNE	FMR	CV	MC	SNE	FMR	CV
0.5	60	43.13	43.12	43.19	43.13	43.20	473.33	29.62	29.98	150.75	0.17	0.18	0.03	0.18
0.5	80	25.51	25.33	25.51	25.47	25.54	382.99	42.39	33.99	133.48	0.12	0.84	0.14	0.28
0.5	100	11.22	10.90	11.15	11.18	11.20	224.72	46.03	31.66	97.02	0.14	2.70	0.40	0.14
0.5	120	3.27	3.05	3.26	3.28	3.29	87.24	25.01	14.42	53.33	0.55	7.48	1.09	0.43
0.5	140	0.87	0.71	0.82	0.87	0.84	28.93	11.68	4.01	23.14	2.97	15.96	2.50	3.95
2	60	43.22	43.19	43.24	43.21	43.25	657.30	35.38	28.87	102.11	0.07	0.12	0.03	0.09
2	80	26.08	25.84	26.04	26.02	26.06	534.99	47.52	29.42	91.18	0.05	0.85	0.09	0.16
2	100	12.84	12.47	12.77	12.80	12.80	325.34	54.82	24.15	69.31	0.31	2.59	0.25	0.03
2	120	4.98	4.71	5.00	4.99	5.01	143.37	31.22	12.18	43.31	0.56	6.09	0.34	0.43
2	140	1.63	1.45	1.63	1.65	1.63	50.79	13.68	3.87	22.93	0.26	11.10	0.50	0.84
5	60	43.26	43.22	43.26	43.25	43.26	797.13	32.46	28.31	70.47	0.00	0.08	0.00	0.02
5	80	26.44	26.24	26.39	26.38	26.39	652.87	57.89	25.50	62.93	0.18	0.60	0.02	0.05
5	100	13.79	13.43	13.71	13.74	13.72	412.26	67.19	17.51	48.80	0.47	2.16	0.08	0.09
5	120	6.09	5.84	6.10	6.10	6.11	202.06	50.99	8.24	31.41	0.33	4.45	0.11	0.19
5	140	2.37	2.19	2.39	2.39	2.39	82.53	18.90	2.74	18.04	1.06	8.41	0.10	0.08
10	60	43.27	43.23	43.25	43.27	43.25	868.54	32.08	28.12	52.85	0.06	0.05	0.01	0.04
10	80	26.58	26.38	26.51	26.51	26.50	716.04	57.68	23.22	47.33	0.28	0.46	0.01	0.04
10	100	14.17	13.84	14.08	14.12	14.09	461.86	72.70	13.72	36.96	0.64	1.76	0.01	0.24
10	120	6.58	6.32	6.59	6.59	6.59	239.04	51.79	5.70	23.94	0.21	4.10	0.05	0.06
10	140	2.76	2.60	2.78	2.79	2.79	106.15	26.82	1.77	14.00	0.97	6.87	0.09	0.21

Table 3.4: Basic Monte Carlo (MC), Small Noise Expansion (SNE), Fast Mean-Reversion (FMR) and Control Variate (CV) price, variance of price and relative error for a European call under the Bates model with different mean-reversion rates (κ) and strikes (K). Relative error is measured against the semi-analytical solution (B).

κ	K	Price				Variance			
		MC	SNE	FMR	CV	MC	SNE	FMR	CV
0.5	60	0.05	0.02	0.04	0.05	0.28	0.02	0.08	0.26
0.5	80	1.07	0.85	0.97	1.04	14.71	7.82	4.38	10.69
0.5	100	5.32	5.04	5.18	5.23	93.33	37.08	17.63	48.79
0.5	120	16.00	15.70	15.79	15.87	248.11	64.98	48.87	107.23
0.5	140	32.16	31.89	31.95	32.00	375.69	83.41	70.81	171.73
2	60	0.07	0.03	0.05	0.07	0.37	0.12	0.04	0.33
2	80	1.49	1.29	1.41	1.46	19.75	17.09	2.41	11.08
2	100	6.78	6.52	6.66	6.66	117.17	57.88	14.77	42.54
2	120	17.60	17.26	17.35	17.43	300.00	74.41	37.27	90.45
2	140	32.80	32.50	32.50	32.59	481.93	97.61	59.16	148.48
5	60	0.08	0.04	0.07	0.08	0.44	0.24	0.01	0.37
5	80	1.83	1.63	1.74	1.79	23.62	25.35	1.61	9.77
5	100	7.73	7.46	7.57	7.60	132.99	71.40	9.90	35.28
5	120	18.75	18.39	18.48	18.56	330.08	91.17	26.72	74.06
5	140	33.64	33.26	33.29	33.41	536.72	96.43	46.91	120.80
10	60	0.09	0.06	0.08	0.09	0.43	0.56	0.01	0.34
10	80	1.98	1.82	1.90	1.94	25.27	33.47	1.19	8.31
10	100	8.20	7.89	8.03	8.05	139.63	73.70	8.04	28.64
10	120	19.36	19.00	19.08	19.16	342.03	97.97	22.62	58.79
10	140	34.22	33.80	33.83	33.97	555.95	98.68	40.70	94.48

Table 3.5: Basic Monte Carlo (MC), Small Noise Expansion (SNE), Fast Mean-Reversion (FMR) and Control Variate (CV) price and variance of price for a down-and-out put under the Bates model with different mean-reversion rates (κ) and strikes (K).

3.5 Greeks

In this section, we show how the Greeks, or derivative hedge ratios, can be calculated by exploiting the variance reduction potential of importance sampling. Broadie and Kaya (2004) compute the Greeks for affine models, including the Heston and the Bates models, using the Likelihood Ratio Method conditioned on the state variables. Glasserman and Liu (2010) also use the Likelihood Ratio Method to calculate the Greeks, but they calculate the likelihood ratio by inverting the Laplace transform of the distribution function using the Abate-Whitt (1992) algorithm. Here, we also use the Likelihood Ratio Method, but one that is obtained via Fourier inversion of the characteristic function. For comparison, we calculated the Greeks according to the method of Benhamou (2002) which is based on the Malliavin calculus and implemented for the Heston model.

We begin with an option price under $\mathbb{Q}$ defined as

$$P(t,x) = \int_0^\infty e^{-r(T-t)}\phi\left(S_T\right)f(x)dx$$

where $\phi\left(S_T\right)$ is the payoff function and $f(x)$ is the risk-neutral probability density function.

Next, consider delta, Δ, the first derivative of the option price with respect to S_0

$$\begin{aligned}\Delta &= \frac{\partial}{\partial S_0}\int_0^\infty e^{-r(T-t)}\phi\left(S_T\right)f(x)dx \\ &= \int_0^\infty e^{-r(T-t)}\phi\left(S_T\right)\frac{\frac{\partial}{\partial S_0}f(x)}{f(x)}f(x)dx \qquad (3.16)\end{aligned}$$

where $\frac{\frac{\partial}{\partial S_0}f(x)}{f(x)}$ is the likelihood ratio. The probability density function is obtained from the characteristic function using the following inversion formula

$$Pr\left(S_T > x\right) = \frac{1}{2} + \frac{1}{\pi}\int_0^\infty Re\left[\frac{\exp\left(-i\omega\ln\left(x\right)\right)\psi_T\left(\omega\right)}{i\omega}\right]d\omega$$

where ψ is the characteristic function. We note that the characteristic functions for both the Heston and the Bates models are available in semi-closed form. Since the total probability mass is 1,

$$F(x) = Pr\left(S_T \leq x\right) = \frac{1}{2} - \frac{1}{\pi}\int_0^\infty Re\left[\frac{\exp\left(-i\omega\ln\left(x\right)\right)\psi_T\left(\omega\right)}{i\omega}\right]d\omega \qquad (3.17)$$

Then the integral in (3.17) can be calculated using the Gauss-Laguerre quadrature. Finally, to get the probability density function we use the following finite difference approximation

$$f(x) \approx \frac{Pr\left(S_T \leq x + \Delta x\right) - Pr\left(S_T \leq x\right)}{\Delta x}$$

To apply importance sampling, note that option price under $\tilde{\mathbb{Q}}$ is

$$P(t,x) = \int_0^\infty e^{-r(T-t)} \phi\left(S_T\right) \frac{f(x)}{g(x)} g(x) dx$$

Delta of an option under the importance sampling measure will be obtained by multiplying the integrand of delta under the risk-neutral measure in (3.16) by $\frac{g(x)}{g(x)}$.

$$\begin{aligned}
\Delta &= \int_0^\infty e^{-r(T-t)} \phi\left(S_T\right) \frac{\frac{\partial f(x)}{\partial S_0}}{f(x)} f(x) \frac{g(x)}{g(x)} dx \\
&= \int_0^\infty e^{-r(T-t)} \phi\left(S_T\right) \frac{\frac{\partial f(x)}{\partial S_0}}{f(x)} \frac{f(x)}{g(x)} g(x) dx
\end{aligned}$$

where $\frac{f(x)}{g(x)}$ is the Radon-Nikodym derivative. Similarly, delta of the down-and-out put option can be calculated as

$$\Delta = \int_0^\infty e^{-r(T-t)} \phi\left(S_T\right) \mathbf{1}\left[\min_{1 \leq i \leq M} S_i > B\right] \frac{\frac{\partial f(x)}{\partial S_0}}{f(x)} \frac{f(x)}{g(x)} g(x) dx$$

where M is the number of time steps. The same procedure as above can be followed to compute the other Greeks such as vega, theta, rho, gamma or vanna.

For comparison, we compute the Greeks using the Malliavin calculus, where the Greek is an expectation of the payoff times a Malliavin weight

$$Greek = \mathbb{E}\left[\phi(S_T) weight\right]$$

The Malliavin weights for delta, Δ and gamma, Γ, of European option are given by

$$\begin{aligned}
weight_\Delta &= \frac{1}{S_o T} \int_0^T \frac{dW_t^S}{\sqrt{v_t}} \\
weight_\Gamma &= -\frac{1}{S_o^2 T} \int_0^T \frac{dW_t^S}{\sqrt{v_t}} + \frac{1}{\left(S_o T\right)^2} \left(\left(\int_0^T \frac{dW_t^S}{\sqrt{v_t}} \right)^2 - \left(\int_0^T \frac{dt}{v_t} \right) \right)
\end{aligned}$$

3.5.1 Heston Greeks

In this section, we present delta, Δ, and gamma, Γ, of the Heston model calculated using Malliavin weights (M), Likelihood Ratio Method (L), Likelihood Ratio Method combined with SNE or FMR importance sampling schemes and that calculated using semi-analytical solution. Relative error is measured against the semi-analytical solution. Tables 3.6 and 3.7 present results for delta and gamma, respectively, of a European option under the Heston model. Tables 3.8 and 3.9 present the results for delta and gamma, respectively, of a barrier option under the Heston model.

It is clear from Table 3.6 that the Likelihood Ratio Method gives significantly lower relative error than the method of calculating the Greeks using the Malliavin weights. The importance sampling further enhances the efficiency with variance reduction for SNE of up to 23 times and FMR of up to 130 times. The superiority of the Likelihood Ratio Method carries on to the gamma calculation presented in Table 3.7. This time, the variance reduction for SNE is up to 30 times. For FMR, the variance reduction is up to 60 times.

For the delta of a down-and-out put, presented in Table 3.8, the variance reduction for SNE is up to 12 times. For FMR, the variance reduction is up to 68 times. For the gamma of a down-and-out put, presented in Table 3.9, the variance reduction for SNE is up to 6 times. For FMR, the variance reduction is up to 70 times.

3.5.2 Bates Greeks

In this section, we present delta, Δ, and gamma, Γ, of the Bates model calculated using Likelihood Ratio Method (L) and Likelihood Ratio Method combined with SNE or FMR importance sampling schemes and that calculated using semi-analytical solution. Relative error is measured against the semi-analytical solution. Tables 3.10 and 3.11 present results for delta and gamma, respectively, of a

κ	K	Delta					Variance				Relative error (%)			
		M	L	SNE	FMR	H	M	L	SNE	FMR	M	L	SNE	FMR
0.5	60	0.82	0.99	0.97	0.97	0.99	240.56	9.42	5.59	5.67	17.3	0.3	1.2	1.3
0.5	80	0.81	0.92	0.92	0.91	0.92	100.28	4.76	1.77	2.24	11.7	0.1	0.0	0.9
0.5	100	0.64	0.69	0.68	0.69	0.70	25.67	2.07	0.39	0.53	7.6	0.7	2.1	1.2
0.5	120	0.27	0.28	0.26	0.29	0.29	3.64	0.80	0.19	0.18	4.6	1.1	8.2	0.5
0.5	140	0.07	0.08	0.07	0.07	0.08	0.57	0.26	0.14	0.22	9.9	2.5	15.7	7.3
2	60	0.99	0.98	0.98	0.98	0.98	10.61	7.08	3.05	3.06	0.9	0.3	0.1	0.1
2	80	0.89	0.89	0.90	0.89	0.89	5.92	4.23	0.91	0.98	0.4	0.1	0.6	0.1
2	100	0.66	0.66	0.66	0.66	0.66	3.11	2.41	0.21	0.24	0.2	0.4	0.6	0.2
2	120	0.36	0.35	0.34	0.36	0.36	1.50	1.25	0.15	0.07	0.6	0.9	3.9	0.2
2	140	0.14	0.14	0.13	0.14	0.14	0.61	0.54	0.10	0.03	0.2	1.2	8.8	0.3
5	60	0.98	0.98	0.98	0.98	0.98	7.49	6.27	2.24	2.25	0.2	0.3	0.2	0.0
5	80	0.87	0.87	0.88	0.87	0.87	4.62	4.02	0.64	0.71	0.0	0.1	0.4	0.0
5	100	0.64	0.64	0.64	0.64	0.64	2.80	2.52	0.13	0.16	0.4	0.2	0.1	0.1
5	120	0.38	0.38	0.37	0.38	0.38	1.59	1.47	0.11	0.04	0.4	0.6	2.0	0.2
5	140	0.18	0.18	0.17	0.18	0.18	0.80	0.74	0.09	0.01	0.1	0.8	4.7	0.4
10	60	0.98	0.98	0.98	0.98	0.98	6.71	5.92	1.93	1.94	0.0	0.6	0.3	0.1
10	80	0.86	0.87	0.87	0.86	0.86	4.29	3.90	0.54	0.60	0.1	0.4	0.4	0.1
10	100	0.63	0.63	0.63	0.63	0.63	2.72	2.54	0.11	0.13	0.5	0.1	0.1	0.1
10	120	0.38	0.38	0.38	0.38	0.38	1.63	1.55	0.11	0.03	0.5	0.1	1.2	0.2
10	140	0.20	0.20	0.19	0.20	0.20	0.89	0.85	0.11	0.01	0.0	0.1	2.9	0.3

Table 3.6: Malliavin (M), Likelihood Ratio Method (L), Likelihood Ratio Method with Small Noise Expansion (SNE) and Likelihood Ratio Method with Fast Mean-Reversion (FMR) Greeks (delta and variance of delta) and relative error for a European call under the Heston model with different mean-reversion rates (κ) and strikes (K). Relative error is measured against the semi-analytical solution (H).

κ	K	Gamma					Variance				Relative error (%)			
		M	L	SNE	FMR	H	M	L	SNE	FMR	M	L	SNE	FMR
0.5	60	-1.4230	-0.0055	0.0000	-0.0003	0.0008	8735.6	0.0513	0.0544	0.0542	168079	753.2	104.3	138.2
0.5	80	-0.7531	-0.0003	0.0028	0.0027	0.0048	2818.4	0.0214	0.0230	0.0231	15758	106.8	42.3	43.2
0.5	100	-0.1790	0.0137	0.0164	0.0162	0.0177	372.94	0.0064	0.0127	0.0076	1113	22.2	7.4	8.3
0.5	120	0.0206	0.0180	0.0185	0.0198	0.0211	7.9288	0.0024	0.0029	0.0026	2.6	14.8	12.4	6.5
0.5	140	0.0194	0.0044	0.0053	0.0056	0.0071	0.4901	0.0013	0.0035	0.0014	174.3	37.4	24.4	20.4
2	60	0.0034	0.0003	0.0008	0.0009	0.0012	0.1358	0.0150	0.0150	0.0150	184.7	78.4	37.7	28.6
2	80	0.0081	0.0056	0.0060	0.0060	0.0063	0.0828	0.0068	0.0070	0.0070	28.0	11.3	5.2	4.3
2	100	0.0161	0.0142	0.0149	0.0149	0.0147	0.0463	0.0023	0.0030	0.0039	9.7	3.4	1.4	1.7
2	120	0.0178	0.0164	0.0164	0.0162	0.0166	0.0229	0.0008	0.0022	0.0018	7.0	1.3	1.3	2.4
2	140	0.0112	0.0103	0.0095	0.0103	0.0103	0.0096	0.0006	0.0003	0.0002	8.1	0.4	7.8	0.3
5	60	0.0016	0.0012	0.0013	0.0014	0.0015	0.0432	0.0096	0.0096	0.0096	6.6	21.4	10.7	4.2
5	80	0.0071	0.0067	0.0069	0.0069	0.0070	0.0301	0.0045	0.0046	0.0047	1.7	3.6	1.4	1.1
5	100	0.0136	0.0132	0.0136	0.0133	0.0134	0.0201	0.0015	0.0017	0.0018	1.8	1.4	1.7	0.4
5	120	0.0148	0.0144	0.0145	0.0143	0.0144	0.0127	0.0004	0.0005	0.0005	3.1	0.2	0.8	0.3
5	140	0.0110	0.0104	0.0102	0.0104	0.0104	0.0075	0.0003	0.0001	0.0001	5.4	0.4	2.4	0.1
10	60	0.0013	0.0016	0.0017	0.0017	0.0016	0.0312	0.0075	0.0076	0.0076	17.7	1.2	2.0	7.7
10	80	0.0070	0.0073	0.0073	0.0073	0.0072	0.0225	0.0036	0.0037	0.0037	2.8	0.4	0.4	0.6
10	100	0.0129	0.0130	0.0132	0.0130	0.0130	0.0156	0.0012	0.0014	0.0014	0.2	0.3	1.8	0.1
10	120	0.0138	0.0137	0.0137	0.0135	0.0135	0.0104	0.0003	0.0004	0.0004	1.7	1.2	1.4	0.1
10	140	0.0105	0.0103	0.0101	0.0101	0.0101	0.0066	0.0002	0.0001	0.0001	3.9	2.0	0.3	0.1

Table 3.7: Malliavin (M), Likelihood Ratio Method (L), Likelihood Ratio Method with Small Noise Expansion (SNE) and Likelihood Ratio Method with Fast Mean-Reversion (FMR) Greeks (gamma and variance of gamma) and relative error for a European call under the Heston model with different mean-reversion rates (κ) and strikes (K). Relative error is measured against the semi-analytical solution (H).

κ	K	Delta				Variance			
		M	L	SNE	FMR	M	L	SNE	FMR
0.5	60	-0.0030	-0.0028	-0.0008	-0.0017	0.0014	0.0008	0.0001	0.0002
0.5	80	-0.0524	-0.0522	-0.0407	-0.0475	0.0847	0.0368	0.0155	0.0109
0.5	100	-0.2567	-0.2564	-0.2428	-0.2484	1.0628	0.2308	0.0844	0.0329
0.5	120	-0.6672	-0.6453	-0.6354	-0.6366	9.3291	0.7078	0.1776	0.1679
0.5	140	-0.8890	-0.8340	-0.8360	-0.8286	40.7111	2.1643	1.3619	1.3908
2	60	-0.0041	-0.0044	-0.0016	-0.0029	0.0017	0.0013	0.0007	0.0001
2	80	-0.0722	-0.0741	-0.0636	-0.0689	0.0808	0.0554	0.0387	0.0043
2	100	-0.2820	-0.2815	-0.2713	-0.2752	0.5015	0.2784	0.0658	0.0222
2	120	-0.5641	-0.5597	-0.5558	-0.5543	1.5981	0.7167	0.1172	0.1183
2	140	-0.7613	-0.7513	-0.7517	-0.7505	3.8517	1.5696	0.6151	0.6255
5	60	-0.0049	-0.0050	-0.0028	-0.0040	0.0020	0.0016	0.0013	0.0000
5	80	-0.0867	-0.0883	-0.0811	-0.0837	0.0856	0.0670	0.0605	0.0025
5	100	-0.2927	-0.2929	-0.2859	-0.2859	0.4428	0.3047	0.0793	0.0183
5	120	-0.5321	-0.5311	-0.5269	-0.5252	1.2288	0.7487	0.1094	0.1253
5	140	-0.7041	-0.7023	-0.6990	-0.7005	2.6741	1.5321	0.5250	0.5635
10	60	-0.0051	-0.0054	-0.0039	-0.0049	0.0019	0.0018	0.0034	0.0000
10	80	-0.0937	-0.0955	-0.0897	-0.0918	0.0853	0.0742	0.0611	0.0018
10	100	-0.2987	-0.3005	-0.2945	-0.2945	0.4153	0.3249	0.0642	0.0195
10	120	-0.5240	-0.5259	-0.5218	-0.5206	1.1042	0.7870	0.1124	0.1435
10	140	-0.6846	-0.6871	-0.6858	-0.6833	2.3388	1.5828	0.5381	0.5937

Table 3.8: Malliavin (M), Likelihood Ratio Method (L), Likelihood Ratio Method with Small Noise Expansion (SNE) and Likelihood Ratio Method with Fast Mean-Reversion (FMR) Greeks (delta and variance of delta) for a down-and-out put under the Heston model with different mean-reversion rates (κ) and strikes (K).

κ	K	Gamma				Variance			
		M	L	SNE	FMR	M	L	SNE	FMR
0.5	60	0.0002	0.0001	0.0000	0.0001	0.0000	0.0000	0.0000	0.0000
0.5	80	0.0031	0.0030	0.0024	0.0028	0.0043	0.0001	0.0001	0.0000
0.5	100	0.0223	0.0146	0.0140	0.0143	0.6736	0.0008	0.0003	0.0001
0.5	120	-0.0355	0.0167	0.0173	0.0169	26.745	0.0039	0.0023	0.0026
0.5	140	-0.2221	0.0015	0.0010	0.0016	189.84	0.0141	0.0128	0.0125
2	60	0.0002	0.0003	0.0001	0.0002	0.0000	0.0000	0.0000	0.0000
2	80	0.0036	0.0039	0.0034	0.0036	0.0007	0.0002	0.0001	0.0000
2	100	0.0104	0.0108	0.0105	0.0105	0.0052	0.0008	0.0001	0.0001
2	120	0.0107	0.0112	0.0114	0.0112	0.0177	0.0023	0.0009	0.0008
2	140	0.0026	0.0034	0.0034	0.0035	0.0417	0.0052	0.0031	0.0031
5	60	0.0003	0.0003	0.0002	0.0002	0.0000	0.0000	0.0000	0.0000
5	80	0.0038	0.0040	0.0038	0.0038	0.0005	0.0002	0.0001	0.0000
5	100	0.0083	0.0087	0.0086	0.0085	0.0025	0.0008	0.0001	0.0001
5	120	0.0073	0.0080	0.0079	0.0078	0.0068	0.0020	0.0007	0.0007
5	140	0.0013	0.0022	0.0020	0.0022	0.0139	0.0041	0.0021	0.0022
10	60	0.0003	0.0003	0.0002	0.0003	0.0000	0.0000	0.0000	0.0000
10	80	0.0038	0.0041	0.0039	0.0039	0.0004	0.0002	0.0001	0.0000
10	100	0.0074	0.0079	0.0079	0.0078	0.0018	0.0008	0.0001	0.0001
10	120	0.0058	0.0066	0.0066	0.0066	0.0048	0.0020	0.0007	0.0007
10	140	0.0002	0.0013	0.0012	0.0012	0.0095	0.0040	0.0020	0.0021

Table 3.9: Malliavin (M), Likelihood Ratio Method (L), Likelihood Ratio Method with Small Noise Expansion (SNE) and Likelihood Ratio Method with Fast Mean-Reversion (FMR) Greeks (gamma and variance of gamma) for a down-and-out put under the Heston model with different mean-reversion rates (κ) and strikes (K).

κ	K	Delta				Variance			Relative error (%)		
		L	SNE	FMR	B	L	SNE	FMR	L	SNE	FMR
0.5	60	1.01	0.98	0.98	0.99	10.77	5.19	5.22	2.1	0.4	0.8
0.5	80	0.93	0.94	0.92	0.92	6.26	5.95	2.23	1.7	2.3	0.3
0.5	100	0.70	0.68	0.69	0.69	3.51	0.53	0.43	1.6	0.6	0.1
0.5	120	0.31	0.28	0.30	0.30	2.01	0.40	0.43	3.8	6.4	0.9
0.5	140	0.10	0.07	0.08	0.09	1.21	0.20	0.49	15.5	14.2	10.9
2	60	0.99	0.98	0.98	0.98	7.01	3.00	3.00	0.6	0.4	0.2
2	80	0.89	0.89	0.89	0.89	4.26	0.95	1.02	0.3	0.7	0.2
2	100	0.66	0.65	0.66	0.66	2.49	0.26	0.26	0.0	0.6	0.2
2	120	0.36	0.35	0.36	0.36	1.33	0.19	0.09	0.2	4.2	0.1
2	140	0.15	0.14	0.15	0.15	0.61	0.12	0.03	1.1	9.8	0.2
5	60	0.99	0.98	0.98	0.98	6.28	2.24	2.25	0.8	0.5	0.3
5	80	0.87	0.87	0.87	0.87	4.08	0.69	0.75	0.5	0.6	0.2
5	100	0.64	0.64	0.64	0.64	2.60	0.17	0.20	0.3	0.2	0.2
5	120	0.38	0.37	0.38	0.38	1.55	0.13	0.05	0.2	2.2	0.1
5	140	0.19	0.18	0.19	0.19	0.81	0.10	0.02	0.7	6.1	0.1
10	60	0.99	0.98	0.98	0.98	5.93	1.95	1.95	1.1	0.5	0.4
10	80	0.87	0.87	0.86	0.86	3.95	0.58	0.64	0.8	0.6	0.3
10	100	0.64	0.63	0.64	0.63	2.61	0.14	0.17	0.6	0.0	0.2
10	120	0.39	0.38	0.39	0.39	1.62	0.12	0.04	0.5	1.7	0.2
10	140	0.21	0.20	0.20	0.20	0.91	0.11	0.01	0.9	4.3	0.2

Table 3.10: Likelihood Ratio Method (L), Likelihood Ratio Method with Small Noise Expansion (SNE) and Likelihood Ratio Method with Fast Mean-Reversion (FMR) Greeks (delta and variance of delta) and relative error for a European call under the Bates model with different mean-reversion rates (κ) and strikes (K). Relative error is measured against the semi-analytical solution (B).

European option under the Bates model. Tables 3.12 and 3.13 present the results for delta and gamma, respectively, of a barrier option under the Bates model.

For the delta of European call, presented in Table 3.10, SNE reduces variance up to 19 times, while FMR gives a variance reduction of up to 76 times. For the gamma of European call, presented in Table 3.11, SNE reduces variance up to 27 times, while FMR gives a variance reduction of up to 50 times.

For the delta of a down-and-out put, presented in Table 3.12, SNE reduces variance up to 12 times, while FMR gives a variance reduction of up to 44 times. For the gamma of a down-and-out put, presented in Table 3.13, SNE reduces variance up to 10 times, while FMR gives a variance reduction of up to 46 times.

κ	K	Gamma				Variance			Relative error (%)		
		L	SNE	FMR	B	L	SNE	FMR	L	SNE	FMR
0.5	60	-0.0066	-0.0022	-0.0019	0.0009	0.1821	0.0422	0.0423	846.2	344.3	320.7
0.5	80	-0.0012	0.0021	0.0027	0.0051	0.1355	0.0215	0.0175	123.3	57.9	46.7
0.5	100	0.0123	0.0168	0.0160	0.0172	0.0993	0.0063	0.0055	28.8	2.3	7.2
0.5	120	0.0166	0.0178	0.0188	0.0205	0.0729	0.0051	0.0087	19.1	13.4	8.6
0.5	140	0.0045	0.0063	0.0053	0.0078	0.0537	0.0026	0.0103	41.8	19.3	31.4
2	60	0.0009	0.0007	0.0008	0.0013	0.0367	0.0142	0.0143	29.3	46.6	39.0
2	80	0.0062	0.0061	0.0061	0.0064	0.0253	0.0067	0.0069	3.2	3.9	4.1
2	100	0.0143	0.0148	0.0145	0.0144	0.0163	0.0027	0.0030	0.1	3.2	0.7
2	120	0.0164	0.0160	0.0162	0.0162	0.0099	0.0010	0.0007	1.3	1.3	0.2
2	140	0.0108	0.0096	0.0104	0.0105	0.0056	0.0008	0.0002	3.1	8.3	0.7
5	60	0.0017	0.0015	0.0016	0.0016	0.0280	0.0095	0.0095	10.0	3.3	1.8
5	80	0.0071	0.0070	0.0070	0.0070	0.0204	0.0045	0.0047	2.6	0.3	0.2
5	100	0.0133	0.0134	0.0131	0.0131	0.0142	0.0016	0.0018	1.5	2.4	0.0
5	120	0.0145	0.0142	0.0141	0.0141	0.0095	0.0004	0.0005	2.2	0.6	0.1
5	140	0.0107	0.0101	0.0104	0.0104	0.0059	0.0003	0.0001	3.3	3.2	0.3
10	60	0.0020	0.0018	0.0019	0.0017	0.0234	0.0075	0.0075	20.5	9.1	13.2
10	80	0.0076	0.0074	0.0074	0.0072	0.0173	0.0036	0.0037	4.5	2.2	2.0
10	100	0.0130	0.0131	0.0128	0.0127	0.0124	0.0013	0.0015	2.5	2.7	0.6
10	120	0.0138	0.0136	0.0134	0.0134	0.0085	0.0003	0.0004	3.1	1.6	0.2
10	140	0.0105	0.0100	0.0101	0.0101	0.0056	0.0002	0.0001	4.0	1.0	0.0

Table 3.11: Likelihood Ratio Method (L), Likelihood Ratio Method with Small Noise Expansion (SNE) and Likelihood Ratio Method with Fast Mean-Reversion (FMR) Greeks (gamma and variance of gamma) and relative error for a European call under the Bates model with different mean-reversion rates (κ) and strikes (K). Relative error is measured against the semi-analytical solution (B).

κ	K	Delta			Variance		
		L	SNE	FMR	L	SNE	FMR
0.5	60	-0.0027	-0.0008	-0.0018	0.0008	0.0001	0.0002
0.5	80	-0.0537	-0.0429	-0.0490	0.0375	0.0196	0.0111
0.5	100	-0.2595	-0.2462	-0.2525	0.2347	0.0910	0.0430
0.5	120	-0.6204	-0.6140	-0.6149	0.7015	0.1981	0.1823
0.5	140	-0.8126	-0.8174	-0.8113	1.9645	1.1706	1.2266
2	60	-0.0040	-0.0015	-0.0029	0.0012	0.0004	0.0001
2	80	-0.0735	-0.0640	-0.0693	0.0532	0.0431	0.0055
2	100	-0.2763	-0.2682	-0.2716	0.2683	0.0927	0.0284
2	120	-0.5428	-0.5411	-0.5395	0.6893	0.1322	0.1322
2	140	-0.7281	-0.7324	-0.7296	1.4952	0.6150	0.6209
5	60	-0.0049	-0.0025	-0.0041	0.0016	0.0008	0.0000
5	80	-0.0869	-0.0788	-0.0830	0.0640	0.0573	0.0038
5	100	-0.2852	-0.2785	-0.2795	0.2901	0.0771	0.0235
5	120	-0.5145	-0.5113	-0.5108	0.7113	0.1208	0.1334
5	140	-0.6794	-0.6805	-0.6798	1.4511	0.5210	0.5536
10	60	-0.0052	-0.0035	-0.0048	0.0017	0.0021	0.0000
10	80	-0.0937	-0.0875	-0.0901	0.0701	0.0634	0.0030
10	100	-0.2922	-0.2856	-0.2870	0.3068	0.0659	0.0245
10	120	-0.5097	-0.5071	-0.5053	0.7425	0.1240	0.1478
10	140	-0.6653	-0.6660	-0.6632	1.4913	0.5268	0.5776

Table 3.12: Likelihood Ratio Method (L), Likelihood Ratio Method with Small Noise Expansion (SNE) and Likelihood Ratio Method with Fast Mean-Reversion (FMR) Greeks (delta and variance of delta) for a down-and-out put under the Bates model with different mean-reversion rates (κ) and strikes (K).

κ	K	Gamma			Variance		
		L	SNE	FMR	L	SNE	FMR
0.5	60	0.0001	0.0000	0.0001	0.0000	0.0000	0.0000
0.5	80	0.0032	0.0025	0.0029	0.0001	0.0001	0.0000
0.5	100	0.0140	0.0134	0.0136	0.0008	0.0003	0.0001
0.5	120	0.0157	0.0162	0.0159	0.0033	0.0019	0.0018
0.5	140	0.0013	0.0013	0.0016	0.0107	0.0091	0.0089
2	60	0.0002	0.0001	0.0002	0.0000	0.0000	0.0000
2	80	0.0038	0.0033	0.0035	0.0002	0.0001	0.0000
2	100	0.0101	0.0099	0.0099	0.0007	0.0002	0.0001
2	120	0.0102	0.0105	0.0103	0.0021	0.0008	0.0008
2	140	0.0027	0.0030	0.0029	0.0047	0.0028	0.0028
5	60	0.0003	0.0001	0.0002	0.0000	0.0000	0.0000
5	80	0.0038	0.0035	0.0037	0.0002	0.0001	0.0000
5	100	0.0081	0.0080	0.0079	0.0007	0.0001	0.0001
5	120	0.0071	0.0072	0.0071	0.0018	0.0007	0.0007
5	140	0.0014	0.0014	0.0014	0.0037	0.0019	0.0020
10	60	0.0003	0.0002	0.0003	0.0000	0.0000	0.0000
10	80	0.0038	0.0037	0.0037	0.0002	0.0001	0.0000
10	100	0.0073	0.0073	0.0072	0.0007	0.0001	0.0001
10	120	0.0058	0.0060	0.0058	0.0017	0.0007	0.0007
10	140	0.0005	0.0005	0.0005	0.0035	0.0018	0.0019

Table 3.13: Likelihood Ratio Method (L), Likelihood Ratio Method with Small Noise Expansion (SNE) and Likelihood Ratio Method with Fast Mean-Reversion (FMR) Greeks (gamma and variance of gamma) for a down-and-out put under the Bates model with different mean-reversion rates (κ) and strikes (K).

3.6 Conclusion

In this chapter, we provide strong evidence that importance sampling reduces the variance of European and barrier options and in most cases it clearly outperforms the control variate. We have also introduced an effective numerical scheme that dramatically improves the speed of both Small Noise Expansion (SNE) and Fast Mean-Reversion (FMR) importance sampling schemes. With an appropriately designed lookup table both SNE and FMR outperform the control variate, not only in terms of variance reduction, but also on speed. Between the two importance sampling schemes, FMR achieves higher variance reduction and exhibits significantly lower relative error, thus FMR is preferred over SNE. Finally, we have shown that Likelihood Ratio Method is an efficient way of obtaining the Greeks, and combining it with importance sampling leads to a significant variance reduction for the Greeks.

Chapter 4

Multi-Level Monte Carlo Simulations with Importance Sampling

4.1 Introduction

In practice, the valuation of multi-asset options typically involves the Monte Carlo simulation. The rate of convergence of this simulation is $\sqrt{n}$ where n is the number of sample paths. Hence, improving the accuracy of the simulation by a factor of 10 requires 100 times as many sample paths. For this reason, variance reduction techniques have become essential. Importance sampling reduces the variance by changing the drift of the simulated sample paths. The extent to which variance reduction is achieved through importance sampling very much depends on the change of drift. Much research effort focuses on how to change the drift to fully exploit the variance reduction potential of importance sampling.

Multi-Level Monte Carlo was first introduced in Giles (2008). It is a Monte Carlo simulation performed on different levels of uniform time discretizations. The main advantage of the Multi-Level Monte Carlo is that, for a given accuracy, it has lower computational cost due to reduced variance compared to the basic Monte Carlo. Here, we show that the variance of Multi-Level Monte Carlo can be further reduced by combining it with importance sampling. Moreover, we demonstrate how this combined scheme can be used to price multi-asset options.

There is relatively little work on variance reduction for multi-asset options in the literature. Barraquand (1995) introduces quadratic resampling and combines it with importance sampling to price European multi-asset options. Avramidis (2002) proposes an algorithm that selects the importance sampling density as a mixture of multivariate Normal densities for best-of Asian and best-of barrier options. Neddermeyer (2011) develops non-parametric importance sampling in conjunction with quasi-random numbers to price basket and best-of options. Giles (2009) uses Multi-Level Monte Carlo to price basket options. The work of Barraquand (1995), Avramidis (2002), Giles (2009), as well as Neddermeyer (2011) is done under the Black-Scholes model.

All three studies, Su and Fu (1999), Arouna (2004) and Caprotti (2008), price basket options using importance sampling with the optimal change of drift obtained by solving an optimization problem. In Su and Fu (1999) the change of drift is based on a stochastic optimization. In Arouna (2004) the change of drift relies on the Robbins-Monro algorithms, whereas in Caprotti (2008) it depends on the least squares minimization. Finally, Pellizzari (1998) suggests the use of control variate based on unconditional and conditional expectations of asset prices as a variance reduction technique for multi-asset options in the Black-Scholes model.

In this chapter, we focus on pricing multi-asset options and incorporating importance sampling in the Multi-Level Monte Carlo simulation. First, we present an application of importance sampling with a stochastic change of drift to multi-asset options. Next, we provide an efficient importance sampling scheme in a Multi-Level Monte Carlo simulation. Then, we combine Multi-Level Monte Carlo with importance sampling to price multi-asset options. In all cases, we explain how the Greeks can be computed in the different simulation schemes using the Likelihood Ratio Method and combine it with importance sampling to reduce the variance of the Greeks.

The remainder of this chapter is organized as follows. In Section 4.2, we present an application of importance sampling with a stochastic change of drift to multi-asset options in the Heston stochastic volatility model and the Bates stochastic volatility model with jumps. We consider basket, best-of, worst-of, spread, absolute, composite and quotient options. In Section 4.3, we extend the Likelihood Ratio Method to multi-asset options and combine it with importance sampling to reduce the variance of the Greeks. In Section 4.4, we derive the optimal change of drift for the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model. In Section 4.5, we apply importance sampling in a Multi-Level Monte Carlo using the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model. We demonstrate that applying importance sampling only on the first level can significantly improve the effective performance of the Multi-Level Monte Carlo. In Section 4.6, we use the Likelihood Ratio Method to estimate the Greeks in a Multi-Level Monte Carlo and again combine it with importance sampling to reduce the variance of the Greeks. In Section 4.7, we combine Multi-Level Monte Carlo with importance sampling for pricing multi-asset options under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model. Finally, Section 4.8 concludes the chapter.

4.2 Importance Sampling for Multi-Asset Options

In this section, we apply importance sampling to price multi-asset options. The dynamics of the multi-asset Heston model under the risk-neutral measure $\mathbb{Q}$ is given by

$$\begin{aligned} dS_{i,t} &= rS_{i,t}dt + \sqrt{v_{i,t}}S_{i,t}dW_{i,t}^{S_i} \\ dv_{i,t} &= \kappa_i\left(\theta_i - v_{i,t}\right)dt + \xi_i\sqrt{v_{i,t}}dW_{i,t}^{v_i} \end{aligned}$$

where $S_{i,t}$ is the i-th stock price, r is the risk-free interest rate, $v_{i,t}$ is the i-th variance, κ_i is the i-th mean-reversion rate, θ_i is the i-th long-term variance, ξ_i is the i-th volatility of volatility and $i = 1, \cdots, n$ denotes the number of underlying assets. The correlation matrix is

$$C = \begin{bmatrix} C_1 & C_2 \\ C_2^\top & C_3 \end{bmatrix} \tag{4.1}$$

where

$$C_1 = \begin{bmatrix} \rho_{1,1} & \cdots & \rho_{1,n} \\ \vdots & \ddots & \vdots \\ \rho_{n,1} & \cdots & \rho_{n,n} \end{bmatrix}$$

is the correlation between the stock price processes,

$$C_2 = \begin{bmatrix} \rho_{1,n+1} & \cdots & \rho_{1,2n} \\ \vdots & \ddots & \vdots \\ \rho_{n,n+1} & \cdots & \rho_{n,2n} \end{bmatrix}$$

is the correlation between the stock price processes and the variance processes, and

$$C_3 = \begin{bmatrix} \rho_{n+1,n+1} & \cdots & \rho_{n+1,2n} \\ \vdots & \ddots & \vdots \\ \rho_{2n,n+1} & \cdots & \rho_{2n,2n} \end{bmatrix}$$

is the correlation between the variance processes.

The difference between the multi-asset Heston model and the multi-asset Bates model is that in the multi-asset Bates model the stock price dynamics under the risk-neutral measure $\mathbb{Q}$ becomes

$$dS_{i,t} = S_{i,t}\left(r - \lambda_i \bar{k}_i\right)dt + S_{i,t}\sqrt{v_{i,t}}dW_{i,t}^{S_i} + S_{i,t}dZ_{i,t}$$

$Z_{i,t}$ is a compound Poisson process with intensity λ_i and log-normal distribution of jump sizes such that if k_i is its jump size then $\ln(1+k_i) \sim \mathcal{N}\left(\ln\left(1+\bar{k}_i\right) - \frac{\delta_i^2}{2}, \delta_i^2\right)$. In the implementation of importance sampling, the jump component is incorporated as a change in volatility of the stock price as shown in Section 3.4, Equation (3.11).

In matrix notation, the dynamics of the multi-asset Heston model is

$$dX_t = b(X_t)\,dt + a(X_t)\,d\eta_t \tag{4.2}$$

where $C = \Sigma\Sigma^\top$ is the correlation matrix in (4.1), η_t is a $2n$-dimensional $\mathbb{Q}$-Brownian motion and

$$X_t = \begin{pmatrix} S_{1,t} \\ \vdots \\ S_{n,t} \\ v_{1,t} \\ \vdots \\ v_{n,t} \end{pmatrix}$$

$$b(x) = \begin{pmatrix} rs_1 \\ \vdots \\ rs_n \\ \kappa_1(\theta_1 - v_1) \\ \vdots \\ \kappa_n(\theta_n - v_n) \end{pmatrix}$$

$$
\begin{aligned}
a(x) &= \Sigma \begin{pmatrix} \sqrt{v_1}s_1 & 0 & \cdots & \cdots & \cdots & 0 \\ 0 & \ddots & \cdots & \cdots & \cdots & \vdots \\ \vdots & \cdots & \sqrt{v_n}s_n & \cdots & \cdots & \vdots \\ \vdots & \cdots & \cdots & \xi_1\sqrt{v_1} & \cdots & \vdots \\ \vdots & \cdots & \cdots & \cdots & \ddots & 0 \\ 0 & \cdots & \cdots & \cdots & 0 & \xi_n\sqrt{v_n} \end{pmatrix} \\
\eta_t &= \begin{pmatrix} Z_t^1 \\ \vdots \\ Z_t^n \\ Z_t^{n+1} \\ \vdots \\ Z_t^{2n} \end{pmatrix}
\end{aligned}
$$

where $Z_t^1, \cdots, Z_t^{2n}$ are independent Brownian motions.

Following Fouque and Tullie (2002), we derive the optimal change of drift for the multi-asset Heston model. First, we introduce the martingale

$$
H_t = \exp\left(\int_0^t h^\top (s, X_s)\, d\eta_s + \frac{1}{2}\int_0^t \|h(s, X_s)\|^2\, ds\right) \tag{4.3}
$$

Next, we define a new probability measure denoted by $\tilde{\mathbb{Q}}$ which is equivalent to $\mathbb{Q}$ by its Radon-Nikodym derivative

$$
\frac{d\tilde{\mathbb{Q}}}{d\mathbb{Q}} = (H_T)^{-1}
$$

By the Girsanov theorem, the process

$$
\tilde{\eta}_t = \eta_t + \int_0^t h(s, X_s)\, ds
$$

is a $2n$-dimensional $\tilde{\mathbb{Q}}$-Brownian motion. Using $\tilde{\eta}_t$, (4.2) and (4.3) can be written as

$$
\begin{aligned}
dX_t &= (b(X_t) - a(X_t)\, h(t, X_t))\, dt + a(X_t)\, d\tilde{\eta}_t \\
H_t &= \exp\left(\int_0^t h^\top (s, X_s)\, d\tilde{\eta}_s + \frac{1}{2}\int_0^t \|h(s, X_s)\|^2\, ds\right)
\end{aligned}
$$

Using the analogous derivation to that presented in Section 3.3.1, the optimal choice of h for which the variance of the Monte Carlo estimator under $\tilde{\mathbb{Q}}$ is minimized is

$$h(t,X_t) = -\frac{1}{P(t,X_t)} a(t,X_t)^\top \nabla P(t,X_t) \tag{4.4}$$

where $P(t,X_t)$ is the option price at time t and $\nabla P(t,X_t)$ is its gradient with respect to x.

This result is also valid for the Bates model with the difference that

$$a(x) = \Sigma \begin{pmatrix} \left(\sqrt{v_1} + \frac{dZ_{1,t}}{dW_{1,t}^{S_1}}\right) s_1 & 0 & \cdots & \cdots & \cdots & 0 \\ 0 & \ddots & \cdots & \cdots & \cdots & \vdots \\ \vdots & \cdots & \left(\sqrt{v_n} + \frac{dZ_{n,t}}{dW_{n,t}^{S_n}}\right) s_n & \cdots & \cdots & \vdots \\ \vdots & \cdots & \cdots & \xi_1\sqrt{v_1} & \cdots & \vdots \\ \vdots & \cdots & \cdots & \cdots & \ddots & 0 \\ 0 & \cdots & \cdots & \cdots & 0 & \xi_n\sqrt{v_n} \end{pmatrix}$$

We refer to Section 3.4 for details.

Equation (4.4) requires the option price and its delta which are not known. Instead, we will use their Black-Scholes equivalents. Under the fast mean-reversion expansion, h for the multi-asset Heston model is given by

$$h_i = -\frac{1}{P_{FMR}} \begin{pmatrix} s_i\sqrt{v_i}\frac{\partial P_{FMR}}{\partial s} \\ 0 \end{pmatrix}$$

where P_{FMR} is the option price under the geometric Brownian motion dynamics with volatility $\sqrt{\sum_{i=1}^n \theta_i^2 - 2\sum_{1\le i<j\le n}^n \rho_{i,j}\sqrt{\theta_i\theta_j}}$.

Similarly, under the fast mean-reversion expansion, h for the multi-asset Bates model is given by

$$h_i = -\frac{1}{P_{FMR}} \begin{pmatrix} s_i\left(\sqrt{v_i} + \frac{dZ}{dW^S}\right)\frac{\partial P_{FMR}}{\partial s} \\ 0 \end{pmatrix}$$

where P_{FMR} is the option price under the geometric Brownian motion dynamics with volatility $\sqrt{\sum_{i=1}^n \theta_i^2 - 2\sum_{1\le i<j\le n}^n \rho_{i,j}\sqrt{\theta_i\theta_j}}$.

4.2.1 Numerical Examples

In this section, we present the numerical results for spread, absolute, composite, quotient, basket, best-of and worst-of options. We compared option prices simulated under the importance sampling using fast mean-reversion expansion (MC+IS) against the basic Monte Carlo (MC). All simulations are performed using the same sequence of pseudo-random numbers. We simulate 10,000 sample paths using a time increment of 0.001. For the numerical examples, we assume that the time to maturity is 1 year. For the Bates model, we assume in addition that the jump intensity is 1 jump per year, standard deviation of the jumps is 2% and the mean jump size is -5%.

4.2.1.1 Spread, Absolute, Composite and Quotient Options

Here, we consider options written on two underlying assets. Spread option depends on the difference between the two underlying assets. Seller of such an option is long correlation which differentiates this option from the majority of multi-asset options that leave the seller short correlation. For example, the payoff of the spread call option with maturity T is given by

$$\max\left(S_1(T) - S_2(T) - K, 0\right)$$

Absolute option is an option written on the absolute value of the difference between the two underlying assets at maturity. The holder of an absolute option benefits from the absolute change in price of the underlying assets. For example, the payoff of the absolute call option with maturity T is given by

$$\max\left(\max\left(S_1(T), S_2(T)\right) - \min\left(S_1(T), S_2(T)\right) - K, 0\right)$$

Composite option is an option on a foreign underlying asset with a strike denominated in the domestic currency. The holder of a composite option faces foreign exchange risk, but benefits from the strike being fixed in the domestic currency. For example, the payoff of the composite call option with maturity T is given by

$$\max\left(S_1(T)S_2(T) - K, 0\right)$$

where $S_2(T)$ is the foreign exchange rate.

Quotient option, also known as ratio option, depends on the ratio of two underlying assets. The holder of a quotient option benefits from the relative change in price of the underlying assets. For example, the payoff of the quotient call option with maturity T is given by

$$\max\left(\frac{S_1(T)}{S_2(T)} - K, 0\right)$$

The parameters used in the numerical examples are displayed in Table 4.1.

i	S	r	v_0	ξ	κ	θ
Spread						
1	30	0.05	0.04	0.4	3	0.09
2	5	0.05	0.09	0.3	0.5	0.25
Absolute						
1	30	0.05	0.04	0.4	3	0.09
2	35	0.05	0.09	0.3	0.5	0.25
Composite and Quotient						
1	30	0.05	0.04	0.4	3	0.09
2	2	0.05	0.09	0.3	0.5	0.25

Table 4.1: Model parameters for multi-asset options based on two underlying assets.

We set the correlation matrix in (4.1) as

$$\begin{bmatrix} 1 & 0.4 & -0.6 & -0.28 \\ 0.4 & 1 & -0.24 & -0.7 \\ -0.6 & -0.24 & 1 & 0.168 \\ -0.28 & -0.7 & 0.168 & 1 \end{bmatrix}$$

Tables 4.2 and 4.3 report the results for basic Monte Carlo (MC) and importance sampling (MC+IS) for the Heston model and the Bates model. In all cases, importance sampling reduces the variance on average 5 to 13 times compared to the basic Monte Carlo.

4.2.1.2 Basket, Best-of and Worst-of Options

Here, we consider options written on three underlying assets. The payoff of a basket option depends on the performance of a basket of underlying assets, each with its own corresponding weight. The weights w_i must satisfy the constraints $0 \leq w_i \leq 1$ for all $i = 1, \cdots, n$ and $\sum_{i=1}^{n} w_i = 1$. For example, the payoff of the basket call option with maturity T is given by

$$\max\left(w_1 S_1(T), \cdots, w_n S_n(T) - K, 0\right)$$

The main advantage of a basket option is that it offers a greater flexibility in the construction of the underlying basket and it is usually cheaper than buying

Moneyness		0.7	0.8	0.9	1	1.1	1.2	1.3
Panel A: Spread								
Heston								
P	MC	10.9116	8.7664	6.8078	5.0821	3.6278	2.4616	1.5897
	MC+IS	10.8962	8.7438	6.7701	5.0316	3.5795	2.4298	1.5730
V	MC	51.5956	47.0730	40.6259	32.9627	25.0319	17.7671	11.7748
	MC+IS	3.4245	3.7729	2.9216	1.7977	1.3673	1.0649	0.7246
Bates								
P	MC	10.9352	8.8015	6.8525	5.1433	3.7056	2.5509	1.6764
	MC+IS	10.9055	8.7716	6.8252	5.1218	3.6918	2.5521	1.6944
V	MC	54.2643	49.5032	42.8706	34.9473	26.7347	19.1973	12.9460
	MC+IS	5.4703	4.7225	4.1768	3.4911	2.5868	1.9515	1.4214
Panel B: Absolute								
Heston								
P	MC	0.8110	0.4772	0.2775	0.1561	0.0860	0.0470	0.0260
	MC+IS	0.8157	0.4738	0.2703	0.1517	0.0837	0.0457	0.0241
V	MC	8.4177	4.9509	2.8182	1.5635	0.8538	0.4603	0.2410
	MC+IS	4.3064	1.9271	0.9153	0.4107	0.1753	0.0765	0.0256
Bates								
P	MC	0.8988	0.5351	0.3091	0.1733	0.0947	0.0509	0.0271
	MC+IS	0.8731	0.5180	0.3015	0.1726	0.0973	0.0530	0.0289
V	MC	9.2814	5.4544	3.0895	1.6995	0.9159	0.4850	0.2541
	MC+IS	3.9397	1.8787	0.8458	0.3651	0.1514	0.0554	0.0207

Table 4.2: Monte Carlo (MC) and Importance Sampling (MC+IS) price (P) and variance of price (V) for spread and absolute options based on two underlying assets under the Heston model and the Bates model.

Moneyness		0.7	0.8	0.9	1	1.1	1.2	1.3
Panel A: Composite								
Heston								
P	MC	32.6826	28.1321	23.9924	20.2822	17.0000	14.1175	11.6163
	MC+IS	32.4569	27.8948	23.7450	20.0351	16.7593	13.9067	11.4605
V	MC	958.7243	888.7356	807.4117	719.3844	629.3578	541.8094	459.6539
	MC+IS	153.3319	155.1247	150.6481	140.7796	127.2177	111.4294	94.7747
Bates								
P	MC	32.7633	28.2312	24.1100	20.4207	17.1358	14.2648	11.7747
	MC+IS	32.6670	28.1174	23.9634	20.2731	17.0256	14.2011	11.7800
V	MC	989.9160	918.5872	835.9709	746.4860	655.8963	567.2028	483.7828
	MC+IS	210.3686	203.9079	189.1749	177.2651	162.4286	145.2400	127.3492
Panel B: Quotient								
Heston								
P	MC	7.2402	5.9534	4.8020	3.8233	3.0250	2.3880	1.8832
	MC+IS	7.2059	5.9192	4.7583	3.7677	2.9560	2.3126	1.8077
V	MC	43.3758	41.5555	38.6297	34.8073	30.5444	26.3034	22.3933
	MC+IS	3.9368	3.6635	3.5537	3.3123	2.9475	2.4886	2.0209
Bates								
P	MC	7.2886	6.0117	4.8691	3.8986	3.0977	2.4557	1.9474
	MC+IS	7.2540	5.9816	4.8495	3.8777	3.0840	2.4433	1.9326
V	MC	44.9052	42.9381	39.8828	35.9234	31.5830	27.2593	23.2442
	MC+IS	6.9516	8.6092	10.8222	11.3142	12.2593	11.4338	8.0210

Table 4.3: Monte Carlo (MC) and Importance Sampling (MC+IS) price (P) and variance of price (V) for composite and quotient options based on two underlying assets under the Heston model and the Bates model.

i	w	S	r	v_0	ξ	κ	θ
Basket							
1	50%	70	0.05	0.04	0.4	3	0.09
2	30%	35	0.05	0.09	0.3	0.5	0.25
3	20%	40	0.05	0.25	0.2	5	0.04
Best-of and Worst-of							
1		30	0.05	0.04	0.4	3	0.09
2		35	0.05	0.09	0.3	0.5	0.25
3		40	0.05	0.25	0.2	5	0.04

Table 4.4: Model parameters for multi-asset options based on three underlying assets.

vanilla options on each of the underlying assets. Basket option is mainly used for diversification purposes.

Best-of option depends on the performance of the best performing asset in a basket. For example, the payoff of the best-of call option with maturity T is given by

$$\max\left(\max\left(S_1(T), \cdots, S_n(T)\right) - K, 0\right)$$

Best-of call option has a higher upside potential compared to a basket call option on the same underlying assets.

Worst-of option depends on the performance of the worst performing asset in a basket. For example, the payoff of the worst-of call option with maturity T is given by

$$\max\left(\min\left(S_1(T), \cdots, S_n(T)\right) - K, 0\right)$$

Worst-of call option has a lower upside potential compared to a basket call option on the same underlying assets.

The parameters used in the numerical examples are displayed in Table 4.4.

We set the correlation matrix in (4.1) as

$$\begin{bmatrix} 1 & 0.4 & 0.2 & -0.6 & -0.28 & -0.1 \\ 0.4 & 1 & 0.5 & -0.24 & -0.7 & -0.25 \\ 0.2 & 0.5 & 1 & 0.0282 & -0.35 & -0.5 \\ -0.6 & -0.24 & 0.0282 & 1 & 0.168 & 0.0294 \\ -0.28 & -0.7 & -0.35 & 0.168 & 1 & 0.175 \\ -0.1 & -0.25 & -0.5 & 0.0294 & 0.175 & 1 \end{bmatrix}$$

Table 4.5 reports the results for basic Monte Carlo (MC) and importance sampling

(MC+IS) for the Heston model and the Bates model. In all cases, importance sampling reduces the variance on average 3 to 5 times compared to the basic Monte Carlo.

4.3 Greeks for Multi-Asset Options

We begin with an option price under $\mathbb{Q}$ defined as

$$P(t,x) = \int_0^\infty \cdots \int_0^\infty e^{-r(T-t)} \phi\left(S_1(T), \cdots, S_n(T)\right) f(x_1, \cdots, x_n) dx_1 \cdots dx_n$$

where $\phi\left(S_1(T), \cdots, S_n(T)\right)$ is the payoff function and $f(x_1, \cdots, x_n)$ is the joint risk-neutral probability density function.

Next, consider, delta Δ, the first derivative of the option price with respect to $S_1(0)$

$$\begin{aligned} \Delta &= \frac{\partial}{\partial S_1(0)} \int_0^\infty \cdots \int_0^\infty B\phi(S_1(T), \cdots, S_n(T)) f(x_1, \cdots, x_n) dx_1 \cdots dx_n \\ &= \int_0^\infty \cdots \int_0^\infty B\phi(S_1(T), \cdots, S_n(T)) LR f(x_1, \cdots, x_n) dx_1 \cdots dx_n \end{aligned}$$

where $B = e^{-r(T-t)}$ is the discount factor, $LR = \frac{\frac{\partial}{\partial S_1(0)} f(x_1, \cdots, x_n)}{f(x_1, \cdots, x_n)}$ is the likelihood ratio and $f(x_1, \cdots, x_n)$ is the joint probability density of the stock price distributions. While $f(x_1, \cdots, x_n)$ is a natural product of the simulation, the partial derivative of $f(x_1, \cdots, x_n)$ with respect to $S_1(0)$ requires further computation. Here, we use an empirically fitted copula to obtain the partial derivative directly. By Sklar's Theorem there exists a copula C such that

$$F\left(x_1, \cdots, x_n\right) = C\left(F_1\left(x_1\right), \cdots, F_n\left(x_n\right)\right) = C\left(u_1, \cdots, u_n\right) \tag{4.5}$$

In Section 3.5, we showed that the cumulative distribution function (CDF) and the probability density function (PDF) for both the Heston model and the Bates model can be obtained as

$$\begin{aligned} F_1\left(x_1\right) &= Pr\left(S_1(T) \le x_1\right) = \frac{1}{2} - \frac{1}{\pi} \int_0^\infty Re\left[\frac{\exp\left(-i\omega \ln\left(x_1\right)\right) \psi_T\left(\omega\right)}{i\omega}\right] d\omega \\ f_1\left(x_1\right) &\approx \frac{F_1\left(x_1 + \Delta x\right) - F_1\left(x_1\right)}{\Delta x} \end{aligned}$$

where ψ is the characteristic function. Taking n^{th} order differentiation of (4.5)

Moneyness		0.7	0.8	0.9	1	1.1	1.2	1.3
Panel A: Basket								
Heston								
P	MC	18.0939	13.4929	9.4321	6.1308	3.6583	2.0060	0.9916
	MC+IS	18.1301	13.5388	9.4302	6.0855	3.6288	1.9740	0.9759
V	MC	131.7882	116.7452	93.7079	66.5853	41.7007	22.8492	11.1104
	MC+IS	22.8241	27.9146	15.9959	11.1225	10.7502	5.6949	1.8306
Bates								
P	MC	18.0972	13.5079	9.4722	6.1935	3.7289	2.0761	1.0521
	MC+IS	18.0389	13.4853	9.4724	6.2015	3.7453	2.0880	1.0719
V	MC	135.8888	120.4728	96.8471	69.1762	43.8538	24.5189	12.2881
	MC+IS	31.0928	26.9779	21.0307	14.3609	8.3187	4.1190	1.7322
Panel B: Best-of								
Heston								
P	MC	20.0806	16.8117	13.6660	10.7524	8.1722	5.9940	4.2450
	MC+IS	20.0837	16.8275	13.6950	10.7276	8.1459	5.9882	4.2398
V	MC	114.2491	112.0612	106.5591	96.5682	82.5957	66.5238	50.5769
	MC+IS	49.6236	40.2004	31.9864	24.6266	19.9598	15.4075	11.2659
Bates								
P	MC	20.2257	16.9608	13.8165	10.8969	8.3088	6.1219	4.3675
	MC+IS	20.1805	16.9499	13.8360	10.8988	8.3327	6.1667	4.4003
V	MC	117.5027	115.1506	109.5496	99.5909	85.5577	69.3061	53.0277
	MC+IS	49.5691	40.9534	33.3357	26.6619	21.6881	16.7958	12.4169
Panel C: Worst-of								
Heston								
P	MC	5.0870	3.1251	1.7429	0.8760	0.3895	0.1544	0.0590
	MC+IS	5.1117	3.1181	1.7279	0.8628	0.3849	0.1524	0.0533
V	MC	32.2608	21.3585	12.1857	5.9854	2.5606	0.9838	0.3423
	MC+IS	13.8514	8.7862	4.7116	2.1009	0.7885	0.2509	0.0701
Bates								
P	MC	5.0152	3.0794	1.7277	0.8753	0.3964	0.1649	0.0615
	MC+IS	5.0174	3.0583	1.6931	0.8489	0.3835	0.1561	0.0580
V	MC	32.4943	21.5437	12.3510	6.1418	2.6990	1.0523	0.3754
	MC+IS	14.3138	8.9745	4.7930	2.1609	0.8268	0.2723	0.0798

Table 4.5: Monte Carlo (MC) and Importance Sampling (MC+IS) price (P) and variance of price (V) for basket, best-of and worst-of options based on three underlying assets under the Heston model and the Bates model.

gives an expression for the joint density.

$$f(x_1, \cdots, x_n) = \prod_{i=1}^{n} f_i(x_i)\, c(u_1, \cdots, u_n) \tag{4.6}$$

where $c(u_1, \cdots, u_n) = \frac{\partial^n C(u_1, \cdots, u_n)}{\partial u_1 \cdots \partial u_n}$ is the copula density. In order to estimate the Greeks, we will use an empirically fitted copula as an approximation of copula in (4.5). The same procedure as above can be followed to compute the other Greeks.

4.3.1 Numerical Examples

In this section, we present delta, Δ, and gamma, Γ, of the multi-asset Heston model and the multi-asset Bates model calculated using Likelihood Ratio Method (MC) and Likelihood Ratio Method combined with importance sampling (MC+IS). We consider basket call option described in Section 4.2.1.2.

The challenge to estimating the Greeks of a basket option is to evaluate the joint density given by (4.6) in a computationally efficient manner. The marginal densities can be easily obtained by creating a lookup table of the cumulative distribution function, which is available in semi-closed form, and then applying the finite difference approximation described in the previous section. However, obtaining copula density is more demanding as there is no readily available semi-closed form result. To overcome this difficulty, we will fit a copula to the data simulated by the multi-asset model. This approach has another advantage that once the copula is fitted, it is possible to create a lookup table of the copula density function and use the finite difference approximation to obtain its derivatives.

To approximate the joint PDF for the Heston model we will use the t-copula with 62 degrees of freedom and correlation between the underlying assets given by

$$\begin{bmatrix} 1 & 0.4 & 0.2 \\ 0.4 & 1 & 0.5 \\ 0.2 & 0.5 & 1 \end{bmatrix} \tag{4.7}$$

To approximate the joint PDF for the Bates model we will use the t-copula with 30 degrees of freedom and the correlation between the underlying assets given by (4.7). For both models the correlation matrix and the number of degrees of freedom were estimated using maximum likelihood.

Tables 4.6 and 4.7 report results for delta and gamma, respectively, of each underlying asset computed by the Likelihood Ratio Method with importance sampling (MC+IS) and without importance sampling (MC) for a basket call

i	Moneyness		0.7	0.8	0.9	1	1.1	1.2	1.3
Heston									
1	D	MC	0.4850	0.4572	0.4070	0.3346	0.2500	0.1678	0.1005
		MC+IS	0.4832	0.4528	0.4005	0.3290	0.2499	0.1702	0.1045
	V	MC	2.9080	2.0993	1.4687	0.9938	0.6445	0.3937	0.2223
		MC+IS	1.9176	1.1636	0.6470	0.3326	0.2544	0.1250	0.0439
2	D	MC	0.2873	0.2671	0.2329	0.1870	0.1366	0.0906	0.0524
		MC+IS	0.2865	0.2673	0.2365	0.1910	0.1352	0.0874	0.0501
	V	MC	7.8665	5.3811	3.4908	2.1292	1.2105	0.6341	0.3048
		MC+IS	6.5965	4.1291	2.3213	1.1894	0.6826	0.2735	0.0868
3	D	MC	0.1886	0.1728	0.1506	0.1219	0.0897	0.0583	0.0340
		MC+IS	0.1932	0.1789	0.1593	0.1315	0.0995	0.0689	0.0430
	V	MC	5.4444	3.6204	2.2803	1.3568	0.7634	0.4083	0.2102
		MC+IS	4.0029	2.3606	1.2534	0.6082	0.2816	0.1127	0.0388
Bates									
1	D	MC	0.4845	0.4567	0.4066	0.3347	0.2512	0.1706	0.1044
		MC+IS	0.4769	0.4489	0.3976	0.3292	0.2488	0.1722	0.1092
	V	MC	2.8944	2.1049	1.4874	1.0193	0.6713	0.4182	0.2421
		MC+IS	1.9074	1.1758	0.6710	0.3565	0.1826	0.0906	0.0463
2	D	MC	0.2933	0.2708	0.2342	0.1862	0.1364	0.0906	0.0529
		MC+IS	0.2812	0.2598	0.2284	0.1803	0.1296	0.0833	0.0477
	V	MC	7.8325	5.3818	3.5098	2.1576	1.2400	0.6602	0.3245
		MC+IS	6.5213	4.1040	2.3413	1.2146	0.5691	0.2377	0.0889
3	D	MC	0.1722	0.1588	0.1390	0.1136	0.0831	0.0543	0.0317
		MC+IS	0.1701	0.1611	0.1442	0.1206	0.0939	0.0655	0.0404
	V	MC	5.2603	3.4916	2.1943	1.3011	0.7289	0.3872	0.1981
		MC+IS	3.8632	2.2887	1.2276	0.6000	0.2648	0.1058	0.0389

Table 4.6: Likelihood Ratio Method (MC) and Likelihood Ratio Method with Importance Sampling (MC+IS) Greeks (delta (D) and variance of delta (V)) for a basket call under the Heston model and the Bates model.

under the Heston model and the Bates model. In both cases, importance sampling reduces the variance of delta and gamma on average by factor of 2.

4.4 Importance Sampling for the Heston Model with Stochastic Interest Rates

Here, we consider models with stochastic volatility and stochastic interest rates. The dynamics of the Heston-Hull-White model and the Heston-Cox-Ingersoll-

i	Moneyness		0.7	0.8	0.9	1	1.1	1.2	1.3
Heston									
1	G	MC	0.0009	0.0025	0.0048	0.0071	0.0084	0.0083	0.0067
		MC+IS	0.0026	0.0039	0.0060	0.0079	0.0089	0.0086	0.0071
	V	MC	0.0378	0.0291	0.0217	0.0157	0.0108	0.0070	0.0042
		MC+IS	0.0262	0.0174	0.0107	0.0061	0.0033	0.0018	0.0009
2	G	MC	0.0115	0.0093	0.0077	0.0065	0.0054	0.0042	0.0027
		MC+IS	0.0169	0.0144	0.0120	0.0100	0.0079	0.0055	0.0036
	V	MC	0.3224	0.2279	0.1536	0.0977	0.0581	0.0322	0.0164
		MC+IS	0.2865	0.1852	0.1098	0.0587	0.0298	0.0127	0.0046
3	G	MC	0.0082	0.0063	0.0047	0.0036	0.0029	0.0024	0.0020
		MC+IS	0.0103	0.0081	0.0062	0.0047	0.0034	0.0025	0.0017
	V	MC	0.1294	0.0883	0.0576	0.0360	0.0218	0.0130	0.0078
		MC+IS	0.0928	0.0554	0.0301	0.0150	0.0079	0.0032	0.0011
Bates									
1	G	MC	0.0028	0.0040	0.0060	0.0080	0.0091	0.0087	0.0071
		MC+IS	0.0040	0.0055	0.0071	0.0089	0.0100	0.0094	0.0078
	V	MC	0.0379	0.0295	0.0223	0.0163	0.0114	0.0076	0.0047
		MC+IS	0.0242	0.0168	0.0105	0.0063	0.0036	0.0020	0.0011
2	G	MC	0.0174	0.0142	0.0115	0.0093	0.0072	0.0053	0.0033
		MC+IS	0.0227	0.0199	0.0163	0.0132	0.0098	0.0072	0.0048
	V	MC	0.3207	0.2294	0.1565	0.1009	0.0608	0.0340	0.0175
		MC+IS	0.2523	0.1679	0.1049	0.0561	0.0279	0.0123	0.0051
3	G	MC	0.0096	0.0071	0.0050	0.0035	0.0025	0.0019	0.0014
		MC+IS	0.0094	0.0075	0.0057	0.0041	0.0032	0.0021	0.0013
	V	MC	0.1184	0.0801	0.0515	0.0315	0.0185	0.0105	0.0059
		MC+IS	0.0781	0.0472	0.0258	0.0129	0.0061	0.0026	0.0009

Table 4.7: Likelihood Ratio Method (MC) and Likelihood Ratio Method with Importance Sampling (MC+IS) Greeks (gamma (G) and variance of gamma (V)) for a basket call under the Heston model and the Bates model.

Ross model under the risk-neutral measure $\mathbb{Q}$ is given by

$$\begin{aligned} dS_t &= r_t S_t dt + \sqrt{v_t} S_t dW_t^S \\ dv_t &= \kappa\left(\bar{v} - v_t\right) dt + \gamma\sqrt{v_t} dW_t^v \\ dr_t &= \lambda\left(\theta_t - r_t\right) dt + \eta r_t^p dW_t^r \end{aligned}$$

where $\left\langle dW_t^S dW_t^v \right\rangle = \rho_{S,v} dt$, $\left\langle dW_t^S dW_t^r \right\rangle = \rho_{S,r} dt$ and $\left\langle dW_t^r dW_t^v \right\rangle = 0$. S_t is the stock price, r_t is the risk-free interest rate, v_t is the variance, κ is the variance mean-reversion rate, $\bar{v}$ is the long-term variance, γ is the volatility of volatility, λ is the interest rate mean-reversion rate, θ_t is the long-term interest rate, η is the volatility of interest rate, $\rho_{S,v}$ is the correlation between stock returns and changes in the variance and $\rho_{S,r}$ is the correlation between stock returns and changes in the interest rate. If $p = 0$, we have the Heston-Hull-White model and if $p = 0.5$, we have the Heston-Cox-Ingersoll-Ross model.

In matrix notation, the model dynamics is

$$dX_t = b\left(X_t\right) dt + a\left(X_t\right) d\eta_t \tag{4.8}$$

where η_t is a 3-dimensional $\mathbb{Q}$-Brownian motion and

$$\begin{aligned} X_t &= \begin{pmatrix} S_t \\ v_t \\ r_t \end{pmatrix} \\ b(x) &= \begin{pmatrix} rs \\ \kappa(\bar{v} - v) \\ \lambda(\theta_t - r) \end{pmatrix} \\ a(x) &= \begin{pmatrix} \sqrt{v}s & 0 & 0 \\ \gamma\sqrt{v}\rho_{S,v} & \gamma\sqrt{v\left(1 - \rho_{S,v}^2\right)} & 0 \\ \eta r^p \rho_{S,r} & \eta r^p \frac{-\rho_{S,v}\rho_{S,r}}{\sqrt{\left(1-\rho_{S,v}^2\right)}} & \eta r^p \sqrt{1 - \left(\rho_{S,r}^2 + \frac{\rho_{S,v}^2 \rho_{S,r}^2}{1-\rho_{S,v}^2}\right)} \end{pmatrix} \\ \eta_t &= \begin{pmatrix} Z_t^1 \\ Z_t^2 \\ Z_t^3 \end{pmatrix} \end{aligned}$$

where Z_t^1, Z_t^2 and Z_t^3 are independent Brownian motions.

Following Fouque and Tullie (2002), we derive the optimal change of drift for the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model. First, we introduce the martingale

$$H_t = \exp\left(\int_0^t h^\top(s, X_s)\, d\eta_s + \frac{1}{2}\int_0^t \|h(s, X_s)\|^2\, ds\right) \tag{4.9}$$

where $h^\top$ denotes the transpose of h. Next, we define a new probability measure denoted by $\tilde{\mathbb{Q}}$ which is equivalent to $\mathbb{Q}$ by its Radon-Nikodym derivative

$$\frac{d\tilde{\mathbb{Q}}}{d\mathbb{Q}} = (H_T)^{-1}$$

By the Girsanov theorem, the process

$$\tilde{\eta}_t = \eta_t + \int_0^t h(s, X_s)\, ds$$

is a 3-dimensional $\tilde{\mathbb{Q}}$-Brownian motion. Using $\tilde{\eta}_t$, (4.8) and (4.9) can be written as

$$\begin{aligned} dX_t &= (b(X_t) - a(X_t)\, h(t, X_t))\, dt + a(X_t)\, d\tilde{\eta}_t \\ H_t &= \exp\left(\int_0^t h^\top(s, X_s) d\tilde{\eta}_s - \frac{1}{2}\int_0^t ||h(s, X_s)||^2 ds\right) \end{aligned}$$

Using the analogous derivation to that presented in Section 3.3.1, the optimal choice of h for which the variance of the Monte Carlo estimator under $\tilde{\mathbb{Q}}$ is minimized is

$$h(t, X_t) = -\frac{1}{P(t, X_t)} a(t, X_t)^\top \nabla P(t, X_t) \tag{4.10}$$

where $P(t, X_t)$ is the option price at time t and $\nabla P(t, X_t)$ is its gradient with respect to x.

From (4.10), under the fast mean-reversion expansion, h is given by

$$h = -\frac{1}{P_{FMR}} \begin{pmatrix} s\sqrt{\bar{v}}\frac{\partial P_{FMR}}{\partial s} \\ 0 \end{pmatrix}$$

where P_{FMR} is the option price under the geometric Brownian motion dynamics with volatility $\sqrt{\bar{v}}$ and interest rate θ.

Strike		60	80	100	120	140
Heston-Hull-White						
Price	HHW	44.0682	26.0077	11.5943	3.7583	0.9221
	MLMC	44.0918	26.0258	11.6100	3.7636	0.9219
	MLMC+IS	44.0879	26.0256	11.6085	3.7648	0.9229
Variance	MLMC	413.4086	379.9792	238.5803	90.0735	23.6667
	MLMC+IS	44.9006	157.6493	124.6953	31.2749	3.4079
Relative error (%)	MLMC	0.05	0.07	0.13	0.14	0.02
	MLMC+IS	0.04	0.07	0.12	0.17	0.08
Heston-Cox-Ingersoll-Ross						
Price	HCIR	44.0686	25.9996	11.5668	3.7296	0.9071
	MLMC	44.0724	26.0067	11.5696	3.7320	0.9037
	MLMC+IS	44.0769	26.0045	11.5693	3.7326	0.9041
Variance	MLMC	413.1371	379.8149	238.3690	90.1313	23.6259
	MLMC+IS	43.8495	159.4084	121.6297	29.2775	3.5179
Relative error (%)	MLMC	0.01	0.03	0.02	0.06	t.37
	MLMC+IS	0.02	0.02	0.02%	0.08	0.33

Table 4.8: Multi-Level Monte Carlo (MLMC) and Importance Sampling (MLMC+IS) price, variance of price and relative error for a European call under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model. Relative error is measured against the semi-analytical solution (HHW/HCIR).

4.5 Multi-Level Monte Carlo with Importance Sampling

As discussed in Section 2.1.6, Multi-Level Monte Carlo is a Monte Carlo simulation with different number of time steps of size $h_l = 2^{-l}T$ on each level $l = 0, 1, \cdots, L$. To illustrate the performance of Multi-Level Monte Carlo with importance sampling, we will use the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model described in the previous section. We compare the performance of the Multi-Level Monte Carlo with importance sampling (MLMC+IS) and without importance sampling (MLMC) for the European call option against the semi-analytical solution (HHW/HCIR). The parameters are set as follows: $\kappa = 2$, $\gamma = 0.06$, $v_0 = \bar{v} = 0.04$, $\lambda = 0.05$, $r_0 = \theta = 0.07$, $\eta = 0.01$, $S_0 = 100$, $\rho_{S,v} = -0.3$, $\rho_{S,r} = 0.2$, $T = 1$. We set $\epsilon = 0.01$ and $L = 8$. We consider five strikes: 60, 80, 100, 120 and 140. Table 4.8 presents the price, variance and relative error of the European call. The relative error is measured against the semi-analytical solution.

Given the nature of the Multi-Level Monte Carlo, we observe that it is possible to use importance sampling on all levels or only some levels. We will refer to the former as full importance sampling. Therefore, as an alternative to the full importance sampling, we will consider importance sampling on the first level only. We note that the first level, where $l = 0$, is the coarsest, because there is only one step at this level with step size T. Moreover, variance at level l decreases as l increases because both P_{l-1} and P_l accurately approximate P as they are obtained using the same Brownian path, which implies that the variance on the first level is the highest.

Figure 4.1 plots the effective performance against strike for Multi-Level Monte Carlo with full importance sampling at every level, Multi-Level Monte Carlo with importance sampling on the first level only and Multi-Level Monte Carlo without importance sampling. Effective performance is defined as the ratio of variance reduction to speed. Speed itself is defined as the ratio of computational time of the Multi-Level Monte Carlo with importance sampling to computational time of the Multi-Level Monte Carlo without importance sampling. The results indicate that the Multi-Level Monte Carlo with full importance sampling is more efficient than Multi-Level Monte Carlo without importance sampling. In addition, Multi-Level Monte Carlo with importance sampling on the first level only is much more efficient than both Multi-Level Monte Carlo without importance sampling and Multi-Level Monte Carlo with full importance sampling. The performance improvement, compared to the Multi-Level Monte Carlo with full importance sampling, comes from two sources. The first one is variance reduction from importance sampling; the second is reduced computational time. This is due to the fact that the number of sample paths at level l which is given by (2.9) depends on the variance at level l. Since importance sampling reduces the variance at the first level, the required number of sample paths at this level is less, compared to the Multi-Level Monte Carlo without importance sampling.

So far, we have demonstrated that importance sampling improves the efficiency of both basic Monte Carlo and Multi-Level Monte Carlo. It has been also shown by Giles (2008) that Multi-Level Monte Carlo is more efficient than basic Monte Carlo. To complete the picture, we compare basic Monte Carlo with Importance Sampling and Multi-Level Monte Carlo with Importance Sampling. Figure 4.2 compares the computational cost associated with the desired accuracy, ϵ, for basic Monte Carlo with Importance Sampling (MC+IS) and Multi-Level

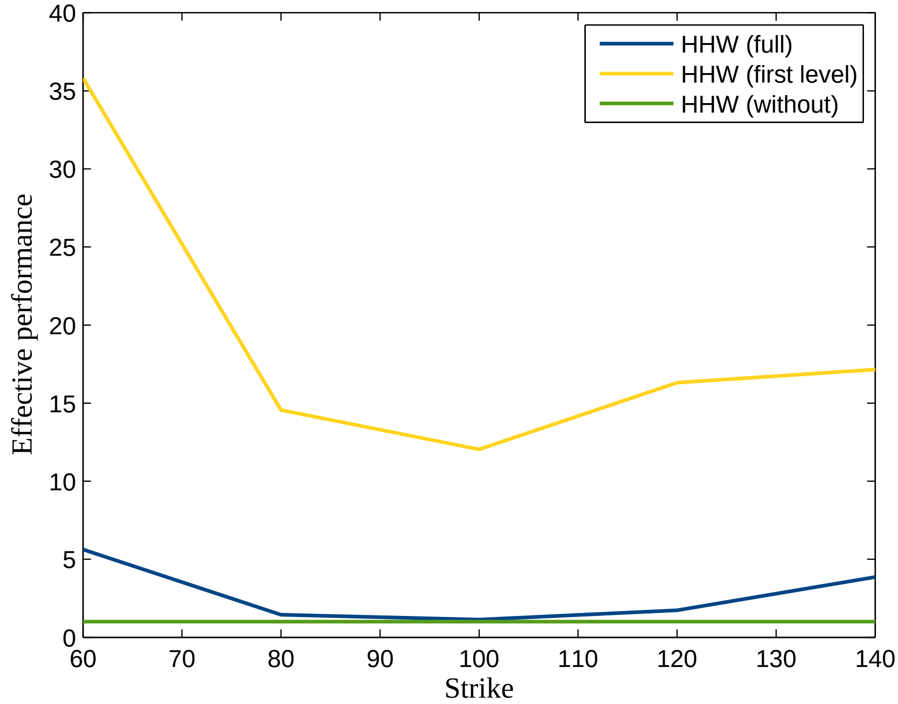

(a) Heston-Hull-White (HHW)

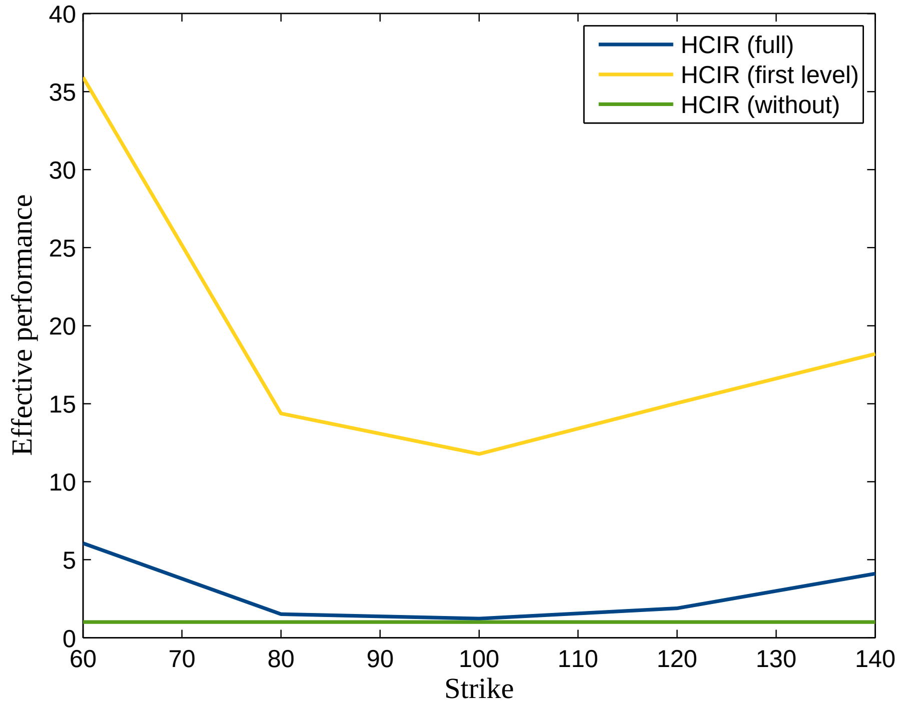

(b) Heston-Cox-Ingersoll-Ross (HCIR)

Figure 4.1: Effective performance of Multi-Level Monte Carlo with full importance sampling, Multi-Level Monte Carlo with importance sampling on the first level only and Multi-Level Monte Carlo without importance sampling for different strikes.

Monte Carlo with Importance Sampling on the first level (MLMC+IS). For Multi-Level Monte Carlo, the computational cost, C, is defined as the total number of time steps on all levels. For each sample path at level $l > 0$, there is one fine path with 2^l time steps and one coarse path with 2^{l-1} time steps. Hence,

$$C = N_0 + \sum_{l=1}^{L} N_l \left(2^l + 2^{l-1}\right)$$

The computational of the basic Monte Carlo is calculated as

$$C^* = \sum_{l=0}^{L} N_l^* 2^l$$

where $N_l^* = 2\epsilon^{-2} Var\left[\phi_l\left(S\right)\right]$ so that the variance of the basic Monte Carlo estimator is also $\frac{1}{2}\epsilon^2$. In Figure 4.2, the dashed line representing MLMC+IS is always below the solid line representing MC+IS. This means that for any level of accuracy Multi-Level Monte Carlo with Importance Sampling on the first level has lower computational cost compared to the basic Monte Carlo with Importance Sampling. Figure 4.2 also shows that MLMC+IS has better convergence properties compared to MC+IS, because for MLMC+IS $\epsilon^2 \cdot C$ is roughly constant, while for MC+IS $\epsilon^2 \cdot C$ is approximately proportional to ϵ^{-1}.

4.6 Greeks for Multi-Level Monte Carlo

It is also possible to use the Likelihood Ratio Method to estimate the Greeks in a Multi-Level Monte Carlo. Let us consider the first derivative of the Multi-Level Monte Carlo estimator (2.6) with respect to S_0.

$$\begin{aligned}
\Delta &= \frac{\partial}{\partial S_0}\mathbb{E}\left[\phi_L\left(S\right)\right] \\
&= \frac{\partial}{\partial S_0}\mathbb{E}\left[\phi_0\left(S\right)\right] + \sum_{l=1}^{L}\frac{\partial}{\partial S_0}\mathbb{E}\left[\phi_l\left(S\right) - \phi_{l-1}\left(S\right)\right] \\
&= \int_0^\infty e^{-r(T-t)}\phi_0\left(s\right)\frac{\frac{\partial}{\partial S_0}f(x)}{f(x)}f(x)dx + \\
&+ \sum_{l=1}^{L}\int_0^\infty e^{-r(T-t)}\left(\left[\phi_l\left(s\right) - \phi_{l-1}\left(s\right)\right]\right)\frac{\frac{\partial}{\partial S_0}f(x)}{f(x)}f(x)dx
\end{aligned}$$

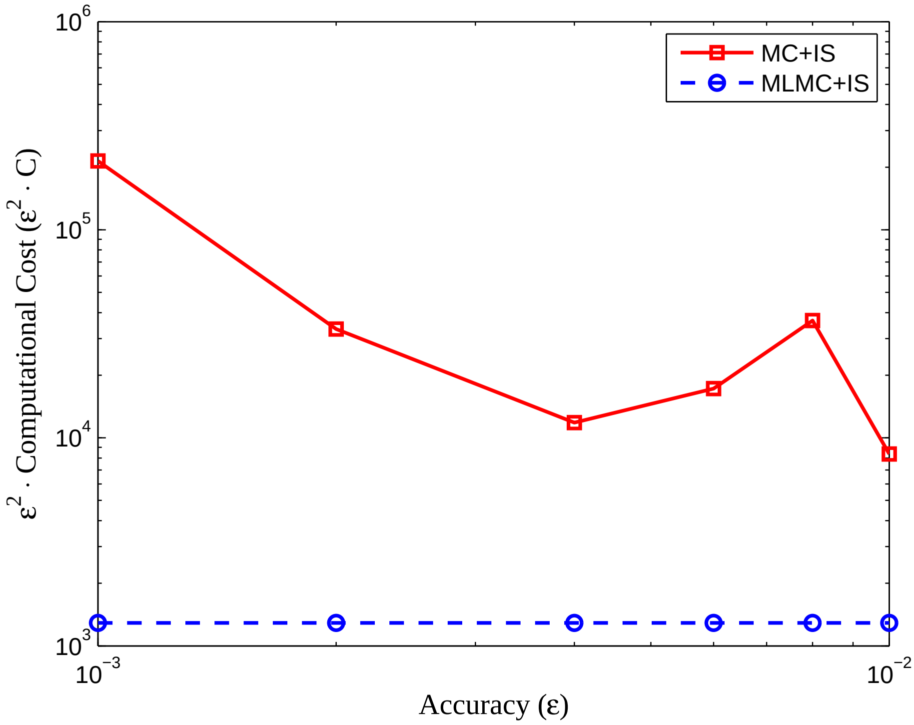

(a) Heston-Hull-White (HHW)

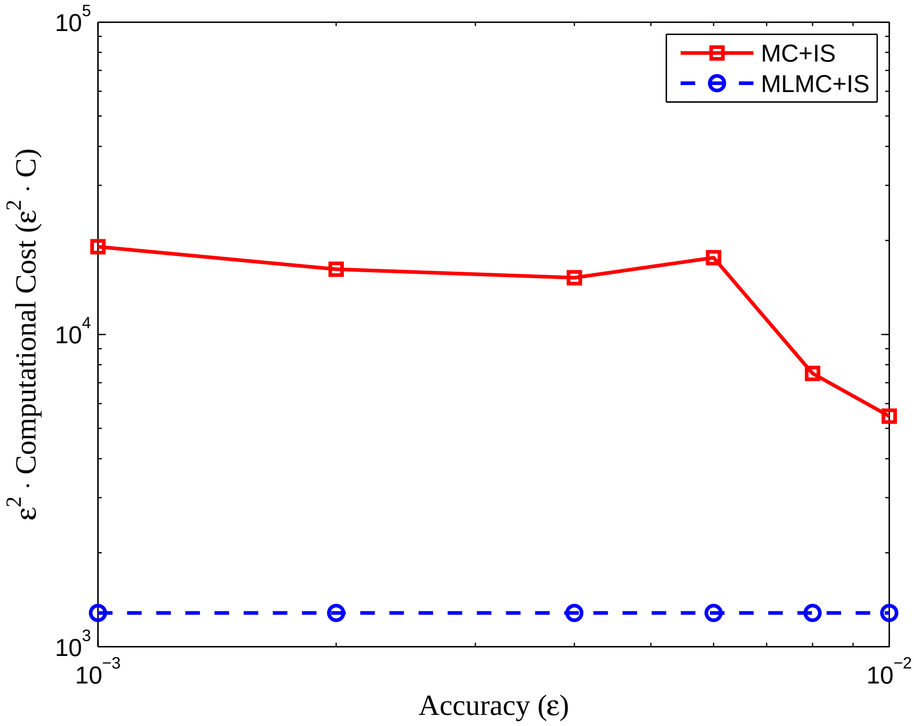

(b) Heston-Cox-Ingersoll-Ross (HCIR)

Figure 4.2: Computational cost associated with the desired accuracy (ϵ) for basic Monte Carlo with Importance Sampling (MC+IS) and Multi-Level Monte Carlo with Importance Sampling on the first level (MLMC+IS).

Strike		60	80	100	120	140
Heston-Hull-White						
D	MLMC	0.9595	0.9187	0.6691	0.3227	0.1068
	MLMC+IS	0.9621	0.9172	0.6690	0.3235	0.1058
V	MLMC	8.5492	4.7461	2.5166	1.1986	0.4324
	MLMC+IS	4.3355	1.1883	0.2358	0.0869	0.0379
G	MLMC	0.0001	0.0042	0.0165	0.0175	0.0092
	MLMC+IS	0.0001	0.0042	0.0164	0.0174	0.0092
V	MLMC	0.0569	0.0370	0.0219	0.0119	0.0055
	MLMC+IS	0.0264	0.0108	0.0035	0.0011	0.0005
Heston-Cox-Ingersoll-Ross						
D	MLMC	0.9570	0.9179	0.6698	0.3231	0.1062
	MLMC+IS	0.9577	0.9195	0.6686	0.3224	0.1057
V	MLMC	8.6794	4.8099	2.5586	1.2152	0.4403
	MLMC+IS	4.4050	1.2070	0.2374	0.0861	0.0350
G	MLMC	0.0001	0.0041	0.0165	0.0176	0.0092
	MLMC+IS	0.0002	0.0042	0.0164	0.0175	0.0091
V	MLMC	0.0591	0.0383	0.0227	0.0123	0.0057
	MLMC+IS	0.0271	0.0112	0.0036	0.0011	0.0005

Table 4.9: Likelihood Ratio Method (MLMC) and Likelihood Ratio Method with Importance Sampling (MLMC+IS) Greeks (delta (D), variance of delta (V), gamma (G) and variance of gamma (V)) for a European call under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model.

where $\frac{\frac{\partial}{\partial S_0} f(x)}{f(x)}$ is the likelihood ratio which can be obtained from the characteristic function as in 3.5.

Table 4.9 reports delta and gamma computed by the Likelihood Ratio Method with importance sampling (MLMC+IS) and without importance sampling (MLMC) for European call option under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model. On average importance sampling reduces the variance of delta 8 times and the variance of gamma 7 times.

4.7 Multi-Level Monte Carlo with Importance Sampling for Multi-Asset Options

Finally, we will use the Multi-Level Monte Carlo to price basket call on three underlying assets. We will use the same parameters as in Section 4.2 and $\lambda = 0.05$, $r_0 = \theta = 0.05$, $\eta = 0.01$, $\rho_{S,v} = -0.3$, $\rho_{S,r} = 0.2$, $\rho_{r,v} = 0$. We set $\epsilon = 0.05$ and $L = 8$. We note that a combination of Multi-Level Monte Carlo and hybrid

Moneyness		0.7	0.8	0.9	1	1.1	1.2	1.3
Heston-Hull-White								
P	MLMC	17.9506	13.2379	9.1427	5.8453	3.3797	1.7681	0.8636
	MLMC+IS	17.9671	13.2424	9.1135	5.7541	3.3249	1.8456	0.9217
V	MLMC	123.9666	114.0966	94.7747	68.4085	43.3949	24.5796	12.3358
	MLMC+IS	41.6504	40.1116	33.2675	22.6223	12.5556	6.4589	2.4066
Heston-Cox-Ingersoll-Ross								
P	MLMC	17.9789	13.2578	9.2080	5.8265	3.3508	1.7697	0.9293
	MLMC+IS	17.9381	13.1990	9.0608	5.7880	3.3921	1.8451	0.9390
V	MLMC	123.7446	114.3640	95.0163	68.7593	43.3255	24.1792	12.3348
	MLMC+IS	41.3532	40.3292	33.2685	22.9515	12.6570	6.1357	2.5758

Table 4.10: Multi-Level Monte Carlo (MLMC) and Importance Sampling (MLMC+IS) price (P) and variance of price (V) for a basket call under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model.

stochastic volatility model such as Heston-Hull-White or Heston-Cox-Ingersoll-Ross is particularly suitable for pricing variable annuities which are in principle long-dated basket put options.

Table 4.10 reports the results for Multi-Level Monte Carlo with importance sampling (MLMC+IS) and without importance sampling (MLMC) for a basket call on three underlying assets under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model. On average importance sampling reduces the variance 3 times compared to the Multi-Level Monte Carlo without importance sampling.

Tables 4.11 and 4.12 report results for delta and gamma, respectively, of each underlying asset computed by the Likelihood Ratio Method with importance sampling (MLMC+IS) and without importance sampling (MLMC) for a basket call under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model. Importance sampling reduces the variance of delta and gamma on average by factor of 3.

4.8 Conclusion

We have presented an application of importance sampling with stochastic change of drift to multi-asset options. We have illustrated the use of importance sampling

i		Moneyness	0.7	0.8	0.9	1	1.1	1.2	1.3
Heston-Hull-White									
1	D	MLMC	0.3564	0.3421	0.2985	0.2464	0.1807	0.1122	0.0639
		MLMC+IS	0.3562	0.3256	0.2956	0.2404	0.1776	0.1201	0.0680
	V	MLMC	2.1769	1.5899	1.1284	0.7723	0.4973	0.3056	0.1876
		MLMC+IS	1.3601	0.8669	0.5247	0.3050	0.1676	0.0910	0.0362
2	D	MLMC	0.5109	0.4553	0.3840	0.3004	0.2275	0.1624	0.1099
		MLMC+IS	0.5040	0.4634	0.3802	0.2979	0.2200	0.1698	0.1174
	V	MLMC	11.6235	8.5318	6.0153	4.2465	2.8229	1.7758	1.1957
		MLMC+IS	8.5810	5.7738	3.7107	2.0393	1.1347	0.6621	0.2517
3	D	MLMC	0.9084	0.8531	0.7083	0.6401	0.4732	0.3304	0.2159
		MLMC+IS	0.8841	0.8444	0.7431	0.6094	0.4847	0.3504	0.2311
	V	MLMC	38.4584	31.8220	25.7602	16.5950	15.8343	11.4764	8.1404
		MLMC+IS	22.4692	13.6315	12.6515	5.0691	3.3503	2.6284	0.6424
Heston-Cox-Ingersoll-Ross									
1	D	MLMC	0.3658	0.3407	0.3069	0.2444	0.1787	0.1158	0.0769
		MLMC+IS	0.3542	0.3327	0.3008	0.2462	0.1831	0.1182	0.0779
	V	MLMC	2.2403	1.6174	1.1399	0.7857	0.5084	0.3250	0.1949
		MLMC+IS	1.3789	0.8831	0.5337	0.3191	0.1751	0.0910	0.0530
2	D	MLMC	0.5209	0.4610	0.3938	0.3160	0.2363	0.1538	0.1066
		MLMC+IS	0.5021	0.4704	0.3833	0.3010	0.2218	0.1693	0.1113
	V	MLMC	11.9429	8.9477	6.1634	4.3160	2.8600	1.6796	1.0906
		MLMC+IS	8.6689	5.9153	3.7445	2.1748	1.1615	0.5881	0.3708
3	D	MLMC	0.9053	0.8298	0.7488	0.6466	0.4987	0.3463	0.2317
		MLMC+IS	0.9361	0.8567	0.7586	0.5951	0.5186	0.3498	0.2279
	V	MLMC	39.3159	29.1621	22.8414	16.5218	10.2947	5.8761	3.4462
		MLMC+IS	19.0239	14.6707	12.4880	3.7593	3.6790	0.7555	0.4583

Table 4.11: Likelihood Ratio Method (MLMC) and Likelihood Ratio Method with Importance Sampling (MLMC+IS) Greeks (delta (D) and variance of delta (V)) for a basket call under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model.

i	Moneyness		0.7	0.8	0.9	1	1.1	1.2	1.3
Heston-Hull-White									
1	G	MLMC	-0.0208	-0.0147	-0.0096	-0.0033	0.0007	0.0022	0.0023
		MLMC+IS	-0.0219	-0.0158	-0.0089	-0.0034	0.0007	0.0026	0.0022
	V	MLMC	0.0190	0.0144	0.0110	0.0079	0.0053	0.0035	0.0025
		MLMC+IS	0.0111	0.0075	0.0047	0.0030	0.0018	0.0010	0.0005
2	G	MLMC	0.0078	0.0145	0.0157	0.0195	0.0190	0.0157	0.0135
		MLMC+IS	0.0073	0.0157	0.0162	0.0173	0.0183	0.0174	0.0135
	V	MLMC	0.4305	0.3349	0.2444	0.1860	0.1268	0.0882	0.0640
		MLMC+IS	0.2910	0.2022	0.1398	0.0784	0.0477	0.0302	0.0117
3	G	MLMC	0.2474	0.2027	0.1676	0.1338	0.0959	0.0687	0.0449
		MLMC+IS	0.2430	0.2081	0.1605	0.1269	0.0970	0.0718	0.0481
	V	MLMC	2.9352	1.9736	1.5644	0.7162	0.4576	0.3758	0.1866
		MLMC+IS	0.7470	1.2402	0.8445	0.1511	0.0870	0.0505	0.0206
Heston-Cox-Ingersoll-Ross									
1	G	MLMC	-0.0205	-0.0156	-0.0091	-0.0038	0.0007	0.0029	0.0030
		MLMC+IS	-0.0212	-0.0154	-0.0086	-0.0032	0.0008	0.0023	0.0037
	V	MLMC	0.0204	0.0148	0.0111	0.0081	0.0057	0.0038	0.0025
		MLMC+IS	0.0112	0.0078	0.0049	0.0031	0.0019	0.0010	0.0010
2	G	MLMC	0.0104	0.0150	0.0166	0.0197	0.0211	0.0163	0.0125
		MLMC+IS	0.0082	0.0165	0.0197	0.0171	0.0178	0.0170	0.0145
	V	MLMC	0.4728	0.3633	0.2544	0.1871	0.1365	0.0780	0.0546
		MLMC+IS	0.2982	0.2181	0.1403	0.0890	0.0502	0.0256	0.0223
3	G	MLMC	0.2529	0.2013	0.1657	0.1364	0.0953	0.0777	0.0465
		MLMC+IS	0.2488	0.2099	0.1667	0.1265	0.1011	0.0683	0.0523
	V	MLMC	2.2489	1.8205	1.3654	1.2104	0.4645	0.3950	0.1975
		MLMC+IS	0.9539	0.4543	0.2702	0.2800	0.0928	0.1580	0.0188

Table 4.12: Likelihood Ratio Method (MLMC) and Likelihood Ratio Method with Importance Sampling (MLMC+IS) Greeks (gamma (G) and variance of gamma (V)) for a basket under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model.

with spread, absolute, composite, quotient, basket, best-of and worst-of options as examples. Based on our results, importance sampling reduces variance of multi-asset options on average by a factor of 3-13.

The chapter has also provided an extension of the Likelihood Ratio Method to multi-asset options and combined it with importance sampling to reduce the variance of the Greeks. Based on our results, importance sampling reduces variance of the Greeks of multi-asset options on average by a factor of 2.

We applied importance sampling in a Multi-Level Monte Carlo and have demonstrated that applying importance sampling on the first level significantly improves its effective performance. For the European option in the Multi-Level Monte Carlo with full importance sampling the effective performance is on average almost 3 times better than that of the Multi-Level Monte Carlo without importance sampling. For the same option in the Multi-Level Monte Carlo with importance sampling on the first level only the effective performance is on average almost 19 times better than that of the Multi-Level Monte Carlo without importance sampling. For the basket option on three underlying assets under the Heston-Hull-White model and the Heston-Cox-Ingersoll-Ross model, importance sampling reduces the variance by factor of 3 compared to the Multi-Level Monte Carlo without importance sampling.

We have also used the Likelihood Ratio Method to estimate the Greeks in a Multi-Level Monte Carlo and combined it with importance sampling to reduce the variance of the Greeks. Based on our results, importance sampling reduces variance of the Greeks on average by a factor of 7-8 for the European option and by a factor of 3 for a basket option.

Appendix A. Characteristic Function for the Heston Model with Stochastic Interest Rates

As shown in Grzelak and Oosterlee (2011), the characteristic function of the Heston-Hull-White model is approximately given by

$$\psi \approx \exp\left(A(u,\tau) + B(u,\tau)\, x_0 + C(u,\tau)\, v_0 + D(u,\tau)\, r_0\right) \tag{4.11}$$

where $x_0 = \log(S_0)$ with

$$\begin{aligned}
A(u,\tau) &= \lambda\theta I_1(\tau) + \kappa\bar{v} I_2(\tau) + \frac{1}{2}\eta^2 I_3(\tau) + \eta\rho_{x,v} I_4(\tau) \\
B(u,\tau) &= iu \\
C(u,\tau) &= \frac{iu-1}{\lambda}\left(1-\exp(-\lambda\tau)\right) \\
D(u,\tau) &= \frac{1-\exp(-D_1\tau)}{\gamma^2\left(1-g\exp(-D_1\tau)\right)}\left(\kappa - \gamma\rho_{x,v} iu - D_1\right)
\end{aligned}$$

and

$$\begin{aligned}
D_1 &= \sqrt{\left(\gamma\rho_{x,v} iu - \kappa\right)^2 - \gamma^2 iu\,(iu-1)} \\
g &= \frac{\kappa - \gamma\rho_{x,v} iu - D_1}{\kappa - \gamma\rho_{x,v} iu - +D_1}
\end{aligned}$$

The integrals $I_1(\tau)$, $I_2(\tau)$, $I_3(\tau)$ and $I_4(\tau)$ can be evaluated analytically or approximated as

$$\begin{aligned}
I_1(\tau) &= \frac{1}{\lambda}(iu-1)\left(\tau + \frac{\exp(-\lambda\tau)-1}{\lambda}\right) \\
I_2(\tau) &= \frac{\tau}{\gamma^2}\left(\kappa - \gamma\rho_{x,v} iu - D_1\right) - \frac{2}{\gamma^2}\log\left(\frac{1-g\exp(-D_1\tau)}{1-g}\right) \\
I_3(\tau) &= \frac{1}{2\lambda^3}(i+u)^2\left(3+\exp(-2\lambda\tau) - 4\exp(-\lambda\tau) - 2\lambda\tau\right) \\
I_4(\tau) &\approx -\left[\frac{b}{c}\left(e^{-ct} - e^{-cT}\right) + a\tau + \frac{a}{\lambda}\left(e^{-\lambda\tau}-1\right) + \frac{b}{c-\lambda}e^{-cT}\left(1-e^{-\tau(\lambda-c)}\right)\right] \\
&\times \frac{1}{\lambda}\left(iu-u^2\right)
\end{aligned}$$

The values a, b and c can be determined by

$$a = \sqrt{\bar{v} - \frac{\gamma^2}{8\kappa}} \tag{4.12}$$

$$b = \sqrt{v_0} - a \tag{4.13}$$

$$c = -\log\left(\frac{\Lambda(1)-a}{b}\right) \tag{4.14}$$

where

$$\Lambda(t) = \sqrt{c(t)(\lambda(t)-1)+c(t)d+\frac{c(t)d}{2(d+\lambda(t))}}$$

and

$$\begin{aligned} c(t) &= \frac{1}{4\kappa}\gamma^2\left(1-e^{-\kappa t}\right) \\ d &= \frac{4\kappa\bar{v}}{\gamma^2} \\ \lambda(t) &= \frac{4\kappa v_0 e^{-\kappa t}}{\gamma^2\left(1-e^{-\kappa t}\right)} \end{aligned}$$

The characteristic function of the Heston-Cox-Ingersoll-Ross model has the same form as (4.11) with

$$\begin{aligned} B(u,\tau) &= iu \\ C(u,\tau) &= \frac{1-\exp(-D_2\tau)}{\gamma^2(1-G_2\exp(-D_2\tau))}(\kappa-\gamma\rho_{x,v}iu-D_2) \\ D(u,\tau) &= \frac{1-\exp(-D_1\tau)}{\eta^2(1-G_1\exp(-D_1\tau))}(\lambda-D_1) \end{aligned}$$

and

$$\begin{aligned} A(u,\tau) &= \int_0^\tau(\kappa\bar{v}C(u,s)+\lambda\theta D(u,s))\,ds \\ &+ \int_0^\tau \rho_{x,r}\eta iu\mathbb{E}\left(\sqrt{v(T-s)}\right)\mathbb{E}\left(\sqrt{r(T-s)}\right)D(u,s)\,ds \end{aligned} \tag{4.15}$$

with

$$\begin{aligned} D_1 &= \sqrt{\lambda^2+2\eta^2(iu-1)} \\ D_2 &= \sqrt{(\gamma\rho_{x,v}iu-\kappa)^2-\gamma^2 iu(iu-1)} \\ G_1 &= \frac{\lambda-D_1}{\lambda+D_1} \\ G_2 &= \frac{\kappa-\gamma\rho_{x,v}iu-D_2}{\kappa-\gamma\rho_{x,v}iu-+D_2} \end{aligned}$$

It is possible to approximate both expectations in (4.15) by deterministic functions of the form $a+b\exp(-c(T-s))$. For the variance process, v, values a, b and c are given by (4.12)-(4.14). For the interest rate process, r, values a, b and

c can be determined by

$$\begin{aligned} a &= \sqrt{\theta - \frac{\eta^2}{8\lambda}} \\ b &= \sqrt{r_0} - a \\ c &= -\log\left(\frac{\Lambda(1) - a}{b}\right) \end{aligned}$$

where

$$\Lambda(t) = \sqrt{c(t)(\lambda(t) - 1) + c(t)\, d + \frac{c(t)\, d}{2(d + \lambda(t))}}$$

and

$$\begin{aligned} c(t) &= \frac{1}{4\lambda}\eta^2\left(1 - e^{-\lambda t}\right) \\ d &= \frac{4\lambda\theta}{\eta^2} \\ \lambda(t) &= \frac{4\lambda r_0 e^{-\lambda t}}{\eta^2\left(1 - e^{-\lambda t}\right)} \end{aligned}$$

Chapter 5

What Does Risk-Neutral Skewness Tell Us about Future Stock Returns?

5.1 Introduction

This chapter examines the relationship between the risk-neutral skewness (RNS) of the returns distribution and expected stock returns. In this chapter, we make two important contributions to the literature. First, we find positive relationship between RNS and future realized stock returns and show that this relationship is driven by the idiosyncratic component of RNS. This is in contrast to Conrad, Dittmar and Ghysels (2013) who find negative relationship between RNS and future stock returns, but they average daily estimates of RNS over a quarter, while we use RNS estimates on the last trading day of each month. Second, we find significant underperformance of low RNS stocks that is conditioned on the stock having high hedging demand, being relatively overvalued and subject to short-sell constraints. Unlike Rehman and Vilkov (2012), who argue that RNS captures stock overvaluation and this feature causes the subsequent underperformance, we find that overvaluation is a conditioning variable; negative RNS stocks that are not overvalued do not subsequently underperform.

We use daily option prices for a large sample of US stocks for the period 1996-2012. We estimate RNS of each stock using the model-free methodology of Bakshi, Kapadia and Madan (2003) and find significant and robust evidence that RNS is positively related to future realized stock returns. A strategy that is long the quintile portfolio with the highest RNS stocks and short the quintile portfolio with the lowest RNS stocks, rebalanced monthly, yields an average return of

61 basis points (t-stat: 2.25) per month, Fama-French (FF) alpha of 49 basis points (t-stat: 2.23) per month and Fama-French-Carhart (FFC) alpha of 55 basis points (t-stat: 2.52) per month. In this case, the return of this strategy is mainly driven by the underperformance of highly negative RNS portfolio. We further decompose RNS into its systematic and idiosyncratic components and find the latter responsible for driving the positive relationship between RNS and future realized stock returns. A strategy that is long the quintile portfolio with the highest risk-neutral idiosyncratic skewness (RNIS) stocks and short the quintile portfolio with the lowest RNIS stocks, rebalanced monthly, yields an average return of 62 basis points (t-stat: 2.40) per month, Fama-French (FF) alpha of 51 basis points (t-stat: 2.39) per month and Fama-French-Carhart (FFC) alpha of 57 basis points (t-stat: 2.68) per month.

Why do stocks exhibiting the most negative RNS subsequently underperform? We find that the significant underperformance of the most negative RNS stocks is mainly driven by stocks within the most negative RNS portfolio that are also overpriced; negative RNS stocks that are not overvalued do not underperform in the subsequent period. We use four proxies for relative overvaluation: market-to-book ratio (Fama and French (1992)), the probability of jackpot return in the following year (Conrad, Kapadia and Xing (2014)), the expected idiosyncratic skewness (Boyer, Mitton and Vorkink (2010)) and the maximum daily stock return in the past month (Bali, Cakici and Whitelaw (2011)). According to these proxies, over-optimistic investors or investors with preference for their lottery-like payoffs have temporarily driven up prices of these stocks, which leads to a subsequent underperformance as the price correction mechanism reverses the overvaluation. Our findings are consistent with the demand-based option pricing model of Garleanu, Pedersen and Poteshman (2009). In their model, option price is a function of end-user demand and the variance of the unhedgable part of the option. We argue that an individual with a large delta exposure to the overpriced stocks, such as a shareholder or a CEO with large amount of stock ownership, will enter into a protective put strategy driving up the demand for out-of-the-money put options thereby making put option price more expensive and lowering the skewness of the risk-neutral distribution. Speculators wanting to gain negative delta exposure, but not able to short-sell the stocks will push up the put option price in the same direction, driving RNS to very low negative values.

Why is negative RNS an informative signal of persistent future stock underperformance and yet the market does not correct this pattern? We find that this significant underperformance of the most negative RNS stocks is mainly driven by stocks within the most negative RNS portfolio that are too costly or too risky to sell short. This is confirmed by four proxies for short-selling constraints: the estimated shorting fee (Boehme, Danielsen and Sorescu (2006)), the relative short interest that captures the demand for short-selling (Asquith, Pathak and Ritter (2005)), firm's size and stock returns' idiosyncratic volatility (Wurgler and Zhuravskaya (2002)). It is conjectured that small firms and high idiosyncratic volatility stocks are harder to short-sell.

Our findings suggest that the predictive ability of highly negative RNS is driven by the demand for put options by those pessimistic investors who are afraid that the stock is currently overvalued, but cannot resort to selling (or short-selling) the stock itself. We confirm that stocks characterized by the highest hedging demand exhibit, on average, significantly more negative RNS relative to the other stocks. In particular, we use four proxies for investor hedging demand. Following Acharya, Lochstoer and Ramadorai (2013), we use Zmijewski's (1984) Z-score to capture firm default risk and the ratio of CEO stock holdings to their base salary to capture the managerial hedging motive. Moreover, we use the ratio of put to all options' trading volume (Taylor, Yadav and Zhang (2009)) and the aggregate open interest across all options for a given maturity (Hong and Yogo (2012)) to directly measure the hedging demand for options.

The rest of the chapter is organized as follows. Section 5.2 reviews related literature. Section 5.3 describes the data. Methodology is described in Section 5.4. Section 5.5 examines the relationship between RNS and future realized stock returns. Section 5.6 investigates the sources of underperformance for the stocks exhibiting the most negative RNS and Section 5.7 concludes.

5.2 Literature Review

Higher moments of the returns distribution are important for portfolio selection (Samuelson (1970); Scott and Horvath (1980); DeMiguel, Plyakha, Uppal and Vilkov (2014)). Asset pricing theory predicts that risk-averse investors like (dislike) positive (negative) skewness and negative (positive) kurtosis of the physical returns distribution. Stocks with more negatively skewed or more leptokurtic returns distribution should earn higher expected returns in order to compensate investors. Harvey and Siddique (2000) develop a model in which coskewness of the physical returns distribution is priced, whereas Brunnermeier, Gollier and Parker (2007), Mitton and Vorkink (2007) and Barberis and Huang (2008) develop models in which investors have preference for positive idiosyncratic skewness of the physical returns distribution.

Higher moments can be modelled either using historical data under the physical probability measure or options data under the risk-neutral probability measure. Although both distributions are related, the latter approach has several advantages. First, higher moments are time-varying and historical higher moments are poor predictors of the future higher moments (Hansis, Schlag and Vilkov (2010)). Second, estimation of historical higher moments requires a long time series of returns which is not always possible (Conrad, Dittmar and Ghysels (2013)). Finally, information embedded in option prices is inherently forward-looking (Bakshi, Cao and Chen (1997)) and reflects investors' risk-neutral expectations about the future evolution of the underlying stock price (Bates (1991)).

There are several other methods of measuring the skewness of the risk-neutral returns distribution in the literature besides the methodology of Bakshi, Kapadia and Madan (2003). Xing, Zhang and Zhao (2009) as well as Bali and Hovakimian (2009) calculate the implied volatility smirk as the difference between the implied volatilities of out-of-the-money put options and at-the-money call options. As Rehman and Vilkov (2012) point out, RNS estimated using the Bakshi, Kapadia and Madan (2003) methodology contains more information than the volatility smirk, because it reflects the skewness of the entire risk-neutral returns distribution. Cremers and Weinbaum (2010) compute the volatility spread as the sum of the differences between the implied volatilities of each pair of call and put options.

The empirical evidence on the relationship between RNS and future stock returns is mixed. Xing, Zhang and Zhao (2010), Rehman and Vilkov (2012) and Bali, Hu and Murray (2013) all find positive relationship between RNS and future stock returns, while Bali and Murray (2013) and Conrad, Dittmar and Ghysels (2013) find negative relationship between RNS and future stock returns. In contrast to Bali and Murray (2013) who consider option portfolio returns, we examine the relationship between RNS and future realized stock returns and find a positive relationship. With respect to Conrad, Dittmar and Ghysels (2013), who average daily estimates of RNS over a quarter and consider portfolios with a holding period of a quarter, we use RNS estimates on the last trading day of each month and consider portfolios with a holding period of a month. Rehman and Vilkov (2012) document that increasing the portfolio holding period as well as increasing the averaging window of RNS decreases the excess return differential between the high RNS and low RNS portfolios. On the other hand, our findings are in agreement with Bali, Hu and Murray (2013), who find a positive relationship between RNS and stock returns, but they use expected stock returns derived from financial analysts' price targets, while we use realized stock returns. The positive relation between RNS and future stock returns is also consistent with the sequential trade model of Easley, O'Hara and Srinivas (1998), where at least some informed investors choose to trade in options before trading in the underlying stocks, and hence option prices carry information that leads stock price movements. In our setup, investors with negative expectations regarding relatively overpriced stocks that are too costly or too risky to sell (short-sell), resort to the option market to hedge against the downside risk, driving RNS to very low negative values. As a result, RNS extracted from option prices contains valuable information that is not incorporated in the current stock prices due to market frictions. As this information is diffused to the stock market over time, prices of relatively overvalued stocks subsequently decrease, giving rise to the positive relationship between RNS and future stock returns.

Our results are also in line with and extend the analysis of Rehman and Vilkov (2012). We show that the positive relationship between RNS and future realized stock returns holds for a longer sample period, including the recent financial crisis. However, unlike Rehman and Vilkov (2012), who argue that RNS proxies for stock overvaluation and causes the subsequent underperformance, we find that it is only the fraction of stocks within the most negative RNS portfolio which are also relatively overpriced that subsequently underperform. In other words, exhibiting highly negative RNS is not a sufficient condition for negative future stock returns, but it is the interplay between highly negative RNS and relative overpricing that causes the underperformance. Moreover, we show that it is the idiosyncratic component of RNS that drives its positive relationship with future stock returns rather than its systematic component.

Previous studies show that RNS is not only related to the future stock returns, but also firm-specific and market characteristics. For example, firm size and market volatility are associated with negative RNS whereas leverage ratio is positively related with RNS (Dennis and Mayhew (2002); Taylor, Yadav and Zhang (2009)), but Toft and Prucyk (1997) report negative relationship of leverage ratio and RNS. Rehman and Vilkov (2012) find that RNS is related to misvaluation. Taylor, Yadav and Zhang (2009) find negative RNS is associated with options market liquidity, whereas Dennis and Mayhew (2002) find negative RNS is associated with stock market liquidity. Taylor, Yadav and Zhang (2009) find firm volatility is negatively related with RNS, whereas Dennis and Mayhew (2002) report the reverse. Han (2008) as well as Taylor, Yadav and Zhang (2009) document negative relationship between the market sentiment and RNS. When investors are more bearish, they might expect the stock price to decrease which could lead to a more negatively skewed future returns distribution. Friesen, Zhang and Zorn (2012) find stocks with greater investor belief differences, small stocks and growth stocks all display more negative RNS. Several researchers find negative relationship between the firm beta and RNS (Dennis and Mayhew (2002); Duan and Wei (2009); Taylor, Yadav and Zhang (2009); Chang, Christoffersen, Jacobs and Vainberg (2012)). Finally, Friesen, Zhang and Zorn (2012) confirm that RNS is explained by firm-specific factors rather than market factors as previously suggested by Dennis and Mayhew (2002) and Taylor, Yadav and Zhang (2009).

5.3 Data Description

We obtain data on option prices from OptionMetrics. We use daily prices of all out-of-the-money options for all stocks from January 1996 to December 2012 for options maturing between 10 and 180 days. We discard options with zero open interest, zero bid price, negative strike, price less than \$0.50 and non-standard settlement. We also remove all duplicate observations on a given day if there is more than one option on the same underlying asset, with the same maturity and strike. Closing option prices are calculated as the averages of the closing bid and ask prices. We use the secondary market 3-month T-Bill rate from the Federal Reserve H.15 release as a proxy for the risk-free interest rate. Data on daily and monthly stock returns and market value are obtained from CRSP. The market value is calculated as the closing share price times the number of shares outstanding at the end of each month. We include only firms that have at least 750 trading days (3 years) of non-missing returns in the past 5 years, non-missing market value and at least 10 trading days of RNS observations in a calendar month. Our sample consists of 128,960 firm-month observations for 3,987 firms from January 1996 to December 2012.

We construct four proxies for relative overvaluation: market-to-book ratio (Fama and French (1992)), the probability of jackpot return in the following year (Conrad, Kapadia and Xing (2014)), the expected idiosyncratic skewness (Boyer, Mitton and Vorkink (2010)) and the maximum daily stock return in the past month (Bali, Cakici and Whitelaw (2011)). High values for each of these proxies have been shown to capture relative overpricing.

The market-to-book ratio is calculated from Compustat data as the market value at the end of the fiscal year divided by the book value of equity. In terms of Compustat data items, the market-to-book ratio is computed as

$$MABA = \frac{PRCC\_F \times CSHO}{CEQ + TXDB}$$

where $MABA$ denotes the market-to-book ratio, $PRCC\_F$ represents share price at the end of fiscal year, $CSHO$ represents the number of shares outstanding, $TXDB$ represents deferred taxes and CEQ represents shareholders' common equity. For the market-to-book ratio, we use December information of the previous year for the period from June of the next year until May of the following year to ensure that the market-to-book ratio is available before portfolio formation.

The jackpot return is defined as a logarithmic return greater than 100% over the next year. We use the estimated regression equation in Conrad, Kapadia and Xing (2012) to calculate the estimate of probability of a firm achieving a jackpot return as

$$\hat{Jackpot} = \frac{\exp(y)}{1 + \exp(y)}$$

where

$$\begin{aligned} y &= -3.29 + 0.06SKEW + 0.18RET12 \\ &- 0.02AGE - 0.25TANG + 0.29SALESGRTH \\ &- 0.43TURN + 0.99STDEV - 0.22SIZE \end{aligned}$$

$SKEW$ denotes the skewness of daily logarithmic returns in the past 3 months, $RET12$ is the logarithmic return in the past year, AGE is the number of years since first appearance on CRSP and $TANG$ is the asset tangibility which is defined as the gross value of property, plant and equipment divided by the total assets. In terms of Compustat data items, the asset tangibility is computed as

$$TANG = \frac{PPEGT}{AT}$$

where $PPEGT$ represents the gross value of property, plant and equipment and AT represents the total assets. $SALESGRTH$ denotes the sales growth defined as a logarithm of sales in year t divided by sales in year $t-1$. In terms of Compustat data items, sales growth is computed as

$$SALESGRTH = \log\left(\frac{SALE_t}{SALE_{t-1}}\right)$$

where $SALE$ represents gross sales. $TURN$ is the average monthly turnover in the past 6 months less the average monthly turnover in the past 18 months, $STDEV$ denotes the volatility of daily returns in the past 3 months and $SIZE$ is the logarithm of market value measured in millions.

Idiosyncratic volatility and idiosyncratic skewness are computed as

$$\begin{aligned} IVol_{i,t} &= \left(\frac{1}{N(t)-1}\sum_{t\in D}\epsilon_{i,t}^2\right)^{1/2} \\ ISkew_{i,t} &= \frac{1}{N(t)-2}\frac{\sum_{t\in D}\epsilon_{i,t}^3}{IVol_{i,t}^3} \end{aligned}$$

where $ISkew$ and $IVol$ denote, respectively, idiosyncratic skewness and idiosyncratic volatility, $\epsilon_{i,t}$ is the regression residual of the Fama and French (1993) three-factor model, D is the set of non-missing daily returns in the past 60 months and $N(t)$ denotes the number of days in D. We require at least 15 observations in the past 60 months to compute idiosyncratic volatility and idiosyncratic skewness.

To estimate the expected idiosyncratic skewness for each firm we follow Boyer Mitton and Vorkink (2010) and first run cross-sectional regression every month

$$\begin{aligned} ISkew_{i,t} &= \gamma_0 + \gamma_1 ISkew_{i,t-60} + \gamma_2 IVol_{i,t-60} + \gamma_3 Mom_{i,t-60} \\ &+ \gamma_4 Turn_{i,t-60} + \gamma_5 NASD_{i,t-60} + \gamma_6 Small_{i,t-60} \\ &+ \gamma_7 Med_{i,t-60} + \Gamma Ind_{i,t-60} + \epsilon_{i,t} \end{aligned} \tag{5.1}$$

$ISkew$ and $IVol$ denote, respectively, idiosyncratic skewness and idiosyncratic volatility computed from daily firm-level residuals of the Fama and French (1993) three-factor model in the past 60 months, Mom denotes the cumulative monthly return in the past 12 months excluding the most recent one and $Turn$ is the average monthly turnover in the past year calculated as the trading volume divided by the number of shares outstanding. Both trading volume and number of shares outstanding are obtained from CRSP. Monthly turnover is averaged in each year with at least 5 valid observations. NASDAQ volume is adjusted for the double counting following Gao and Ritter (2010): NASDAQ volume is divided by 2 for the period from 1983 to January 2001, by 1.8 for the rest of 2001, by 1.6 for 2002-2003 and is unchanged from January 2004 to December 2012. $NASD$ is 1 if the firm is listed on NASDAQ and 0 otherwise, $Small$ is 1 if the firm is in the bottom three size deciles and 0 otherwise, Med is 1 if the firm is in one of the size deciles between fourth and seventh and 0 otherwise and Ind is 1 if the firm belongs to a certain industry and 0 otherwise. We define 30 industries as in Fama and French (1997). Finally, we use the fitted part of (5.1) as the estimate of the expected idiosyncratic skewness.

We also construct four measures of short-sale constraints: estimated shorting fee (Boehme Danielsen and Sorescu (2006)), relative short interest (Asquith Pathak and Ritter (2005)), idiosyncratic volatility (Wurgler and Zhuravskaya (2002)) and size. High values for the first three proxies and small firm size typically reflect severe short-selling constraints.

Shorting fee represents the opportunity cost incurred by the short-seller. We use the estimated regression equation in Boehme, Danielsen and Sorescu (2006) to calculate the estimate of the shorting fee as

$$\begin{aligned} \hat{Fee} &= 0.07834 + 0.05438VRSI - 0.00664VRSI^2 \\ &+ 0.000382VRSI^3 - 0.5908Option + 0.2587OptionVRSI \\ &- 0.02713OptionVRSI^2 + 0.0007583OptionVRSI^3 \end{aligned}$$

where RSI is the relative short interest, $VRSI$ is the vicile rank of RSI (1 if the firm is in the bottom 5% in terms of RSI, 2 if the firm is between the 5th and 10th percentile, etc.) and *Option* is 1 if there is non-zero trading volume for the option in the month and 0 otherwise. We obtain the short interest data from Compustat. Relative short interest is defined as the outstanding shorts reported by NYSE and NASDAQ divided by the number of shares outstanding and proxies the supply of shares available to be shorted.

Finally, we construct four variables that proxy for the hedging demand: put/all option volume ratio (Taylor, Yadav and Zhang (2009)), open interest (Hong and Yogo (2012)), CEO stock holdings to base salary ratio (Acharya, Lochstoer and Ramadorai (2013)) and the Zmijewski (1984) score. Higher hedging demand drives up put option volume and open interest. CEO stock holdings to base salary ratio captures background risk. Guiso and Paiella (2008) find that background risk increases risk aversion, whereas Acharya, Lochstoer and Ramadorai (2013) argue that managerial risk-aversion increases with the probability of default and that managerial risk-aversion to default drives the hedging demand. Zmijewski (1984) score measures the probability of default. The higher the Zmijewski (1984) score, the higher the probability of default.

To calculate the put/all option volume ratio, we sum option volume each day across all available strikes for a given maturity. Similarly, open interest is calculated as the sum of open interest of all put and call options each day for a given maturity. The options that are used to compute the put/all option volume ratio and open interest have the same maturity as the options used to estimate RNS.

The CEO stock holdings to base salary ratio is calculated from Execucomp and Compustat data as the number of shares owned by the executive times the share price divided by the salary of the executive. In terms of Execucomp and Compustat data items, the CEO stock holdings to base salary ratio is computed as

$$CEO = \frac{PRCC\_F \times SHROWN\_EXCL\_OPTS}{SALARY}$$

where CEO denotes CEO stock holdings to base salary ratio, $PRCC\_F$ represents share price at the end of fiscal year, $SHROWN\_EXCL\_OPTS$ represents the number of shares owned by the executive, excluding options and $SALARY$ represents base salary of the executive.

In terms of Compustat data items the Zmijewski (1984) score is computed as

$$Z = -4.3 - 4.5\frac{NI}{AT} + 5.7\frac{LT}{AT} - 0.004\frac{ACT}{LCT}$$

where Z denotes the Zmijewski (1984) score, NI represents net income, AT represents total assets, LT represents total liabilities, ACT represents total current assets and LCT represents total current liabilities.

5.4 Methodology

Risk-neutral moments are estimated using the results in Bakshi and Madan (2000) and Bakshi, Kapadia and Madan (2003). Bakshi and Madan (2000) prove that any payoff function $H(S)$ that is twice continuously differentiable can be spanned by a portfolio of zero-coupon bond, stock, out-of-the money put and call options as follows

$$\begin{aligned} \mathbb{E}\left[\exp(-r\tau) H(S)\right] &= \left(H\left(\bar{S}\right) - \bar{S}H_S\left(\bar{S}\right)\right)\exp(-r\tau) + H_S\left(\bar{S}\right)S \\ &+ \int_{\bar{S}}^{\infty} H_{ss}(K)\, C(t,\tau;K)\, dK + \int_0^{\bar{S}} H_{ss}(K)\, P(t,\tau;K)\, dK \end{aligned} \tag{5.2}$$

where $H_S\left(\bar{S}\right)$ denotes the first-order derivative of the payoff function with respect to S evaluated at $\bar{S}$ and $H_{ss}(K)$ denotes the second-order derivative of the payoff function with respect to S evaluated at K.

Bakshi, Kapadia and Madan (2003) define the variance, the cubic and the quadratic contract as

$$H(S) = \begin{cases} R(t,\tau)^2 \\ R(t,\tau)^3 \\ R(t,\tau)^4 \end{cases}$$

where $R(t,\tau) = \log(S(t+\tau)) - \log(S(t))$ is the τ-period stock return.

Furthermore, Bakshi, Kapadia and Madan (2003) show that risk-neutral variance, skewness and kurtosis are given by

$$\begin{aligned}
RNV_t(\tau) &= \exp(r\tau) V_t(\tau) - \mu_t(\tau)^2 \\
RNS_t(\tau) &= \frac{\exp(r\tau)(W_t(\tau) - 3\mu_t(\tau) V_t(\tau)) + 2\mu_t(\tau)^3}{\left[\exp(r\tau) V_t(\tau) - \mu_t(\tau)^2\right]^{3/2}} \\
RNK_t(\tau) &= \frac{\exp(r\tau)\left(X_t(\tau) - 4\mu_t(\tau) W_t(\tau) + 6\mu_t(\tau)^2 V_t(\tau)\right) - 3\mu_t(\tau)^4}{\left[\exp(r\tau) V_t(\tau) - \mu_t(\tau)^2\right]^2}
\end{aligned}$$

where $V_t(\tau)$, $W_t(\tau)$ and $X_t(\tau)$ can be calculated based on (5.2) as follows

$$\begin{aligned}
V_t(\tau) &= \int_{S_t}^{\infty} \frac{2\left(1 - \log\left(\frac{K}{S_t}\right)\right)}{K^2} C_t(\tau; K)\, dK \qquad (5.3) \\
&+ \int_0^{S_t} \frac{2\left(1 + \log\left(\frac{S_t}{K}\right)\right)}{K^2} P_t(\tau; K)\, dK \\
W_t(\tau) &= \int_{S_t}^{\infty} \frac{6\log\left(\frac{K}{S_t}\right) - 3\left(\log\left(\frac{K}{S_t}\right)\right)^2}{K^2} C_t(\tau; K)\, dK \qquad (5.4) \\
&- \int_0^{S_t} \frac{6\log\left(\frac{S_t}{K}\right) + 3\left(\log\left(\frac{S_t}{K}\right)\right)^2}{K^2} P_t(\tau; K)\, dK \\
X_t(\tau) &= \int_{S_t}^{\infty} \frac{12\left(\log\left(\frac{K}{S_t}\right)\right)^2 - 4\left(\log\left(\frac{K}{S_t}\right)\right)^3}{K^2} C_t(\tau; K)\, dK \qquad (5.5) \\
&+ \int_0^{S_t} \frac{12\left(\log\left(\frac{S_t}{K}\right)\right)^2 + 4\left(\log\left(\frac{S_t}{K}\right)\right)^3}{K^2} P_t(\tau; K)\, dK
\end{aligned}$$

and

$$\mu_t(\tau) = \exp(r\tau) - 1 - \frac{\exp(r\tau)}{2} V_t(\tau) - \frac{\exp(r\tau)}{6} W_t(\tau) - \frac{\exp(r\tau)}{24} X_t(\tau)$$

Integrals given by equations (5.3)-(5.5) require a continuum of option prices. As traded options are available at discrete strikes, we use Simpson's rule to approximate these integrals between moneyness of 1/3 and 3. Moreover, we require there are at least two out-of-the-money put options and two out-of-the-money call options to approximate integrals in (5.3)-(5.5). For a given maturity, we interpolate the implied volatility using a piecewise cubic Hermite interpolating polynomial between the lowest and the highest available moneyness separately for put and call options. We extrapolate the implied volatility outside the lowest and the highest available moneyness using implied volatility at the boundaries. We then use the Black-Scholes formula to get the option prices from the implied volatilities and use these option prices to compute the risk-neutral moments. Finally, we use the risk-neutral moments on the last trading day of the month with the shortest maturity if on that day there is more than one maturity available.

In Figure 5.1, we plot the 1st, 10th, 25th, 50th, 75th, 90th and 99th percentiles of the monthly cross-sectional distribution of RNS. All percentiles are time-varying, but the 1st and 99th percentiles show the greatest variation, especially in 1998, 2001, 2008 and 2010-2012. The median RNS is negative, which indicates that the risk-neutral returns distribution is left-skewed. The average maturity of the options used is 86.69 days.

To derive the risk-neutral comoments, Bakshi, Kapadia and Madan (2003) use the following one-factor model of stock returns

$$r_{i,t} = a_i + b_i r_{m,t} + e_{i,t} \tag{5.6}$$

where $r_{i,t}$ is the individual stock return, $r_{m,t}$ is the market return and $e_{i,t}$ is the idiosyncratic risk with zero mean and is independent of $r_{m,t}$. Equation (5.6) is also well-defined under the risk-neutral probability measure. Thus, the risk-neutral comoments are given by

$$\begin{aligned}
RNCV_i(t,\tau) &= b_i RNV_m(t,\tau) \\
RNCS_i(t,\tau) &= b_i RNS_m(t,\tau)\frac{RNV_m(t,\tau)}{\sqrt{RNV_i(t,\tau)}} \\
RNCK_i(t,\tau) &= b_i \frac{RNK_m(t,\tau)}{RNV_i(t,\tau)}
\end{aligned}$$

where

$$b_i = \frac{S_{i,t}}{C_{i,t}} N\left(\frac{\log\left(\frac{S_{i,t}}{K_i}\right) + \left(r - d_i + 0.5\sigma_i^2\right)\tau}{\sigma_i\sqrt{\tau}}\right)\beta_i$$

$N(x)$ denotes the standard normal cumulative distribution function, d_i is the

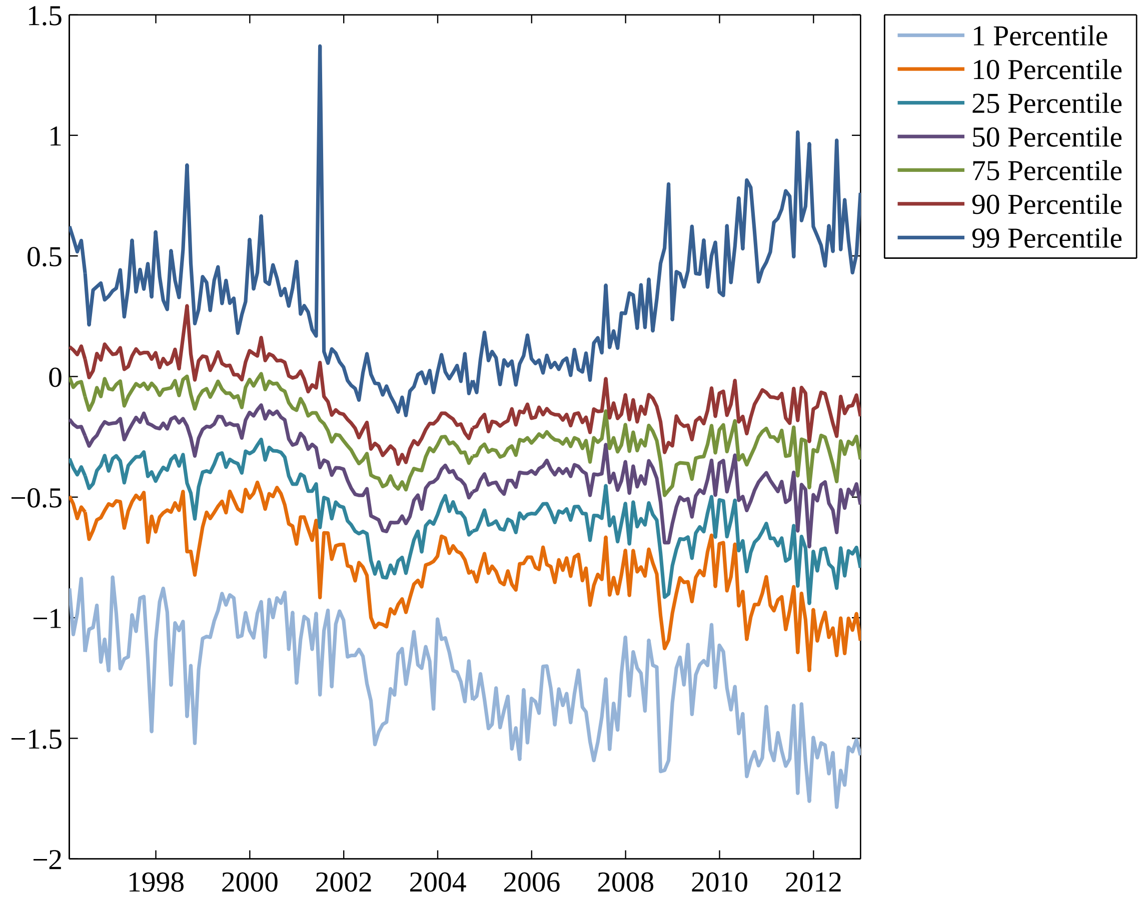

Figure 5.1: Cross-sectional distribution of the risk-neutral skewness.

dividend yield, σ_i is the standard deviation of daily stock returns and β_i is the slope coefficient from regressing daily stock returns on daily S&P 500 returns. We measure the dividend yield using the put-call parity for put and call options with the same strike closest to the stock price. Next, we require at least 120 days of non-missing return data in the past year to estimate β_i and σ_i. Finally, $\frac{S_{i,t}}{C_{i,t}}$ and $\frac{S_{i,t}}{K_i}$ are calculated as the average ratios across call options.

To calculate the risk-neutral idiosyncratic moments, we follow Conrad, Dittmar and Ghysels (2013) and regress the risk-neutral moments on the risk-neutral co-moments for each firm.

$$
\begin{aligned}
RNV_{i,t} &= \kappa_{0i}^V + \kappa_{1i}^V RNCV_{i,t} + \zeta_{i,t}^V \\
RNS_{i,t} &= \kappa_{0i}^S + \kappa_{1i}^S RNCS_{i,t} + \zeta_{i,t}^S \\
RNK_{i,t} &= \kappa_{0i}^K + \kappa_{1i}^K RNCK_{i,t} + \zeta_{i,t}^K
\end{aligned}
$$

We use the sum of the intercept and the residual value as the risk-neutral idiosyncratic moment. For both risk-neutral comoments and idiosyncratic moments, we use observations on the last trading day of the month with the shortest maturity if on that day there is more than one maturity available.

5.5 The Relationship between Risk-Neutral Skewness and Future Stock Returns

5.5.1 Risk-Neutral Skewness Portfolio Sorts

In this section, we examine the relationship between RNS extracted from daily option prices for individual stocks and their future realized returns. At the end of each month, firms are ranked into terciles (quintiles) based on their RNS in the past month. Each tercile (quintile) contains 33⅓% (20%) of the sample. We then form three (five) equally weighted portfolios on the basis of RNS. Portfolios are formed every month from February 1996 to December 2012. Table 5.1 reports the average portfolio returns as well as their Fama-French (α_{FF}) and Fama-French-Carhart (α_{FFC}) alphas estimated from the corresponding 3-factor and 4-factor models, during the period 1996-2012.

The return differential between the highest RNS and lowest RNS quintile portfolios is on average 61 basis points per month (t-stat: 2.25), α_{FF} of 49 basis

Panel A: Terciles														
Tercile	RNS	Mean return	α_{FF}	α_{FFC}	β_{MKT}	β_{SMB}	β_{HML}	β_{MOM}	R^2	RNV	RNK	Ln MV	Std. dev. of ln MV	N
1	-0.7166	0.48	-0.33	-0.32	1.10	0.30	0.01	-0.01	0.93	0.4310	3.4930	22.5612	1.5511	212
			(-3.05)	(-2.97)	(43.68)	(9.24)	(0.25)	(-0.45)						
2	-0.3877	0.74	-0.14	-0.13	1.20	0.44	-0.06	-0.02	0.93	0.4799	3.1073	22.1134	1.4156	212
			(-1.15)	(-1.03)	(42.97)	(12.22)	(-1.64)	(-0.89)						
3	-0.1316	1.01	0.09	0.15	1.24	0.61	-0.23	-0.09	0.90	0.5461	3.0282	21.5752	1.3766	212
			(0.52)	(0.85)	(30.58)	(11.78)	(-4.23)	(-2.61)						
3-1	0.5850	0.52	0.42	0.47	0.14	0.31	-0.24	-0.08	0.39	0.1151	-0.4649	-0.9861	-0.1745	
t(3-1)	81.80	2.34	2.36	2.65	3.52	6.00	-4.33	-2.30		38.52	-26.14	-29.50	-14.78	
Panel B: Quintiles														
Quintile	RNS	Mean return	α_{FF}	α_{FFC}	β_{MKT}	β_{SMB}	β_{HML}	β_{MOM}	R^2	RNV	RNK	Ln MV	Std. dev. of ln MV	N
1	-0.8268	0.46	-0.32	-0.32	1.07	0.28	-0.02	0.00	0.90	0.4224	3.6497	22.6520	1.5817	127
			(-2.59)	(-2.57)	(36.71)	(7.52)	(-0.39)	(0.03)						
2	-0.5249	0.56	-0.29	-0.29	1.16	0.34	0.03	-0.01	0.93	0.4487	3.2313	22.3872	1.4742	127
			(-2.53)	(-2.45)	(42.93)	(9.80)	(0.83)	(-0.45)						
3	-0.3866	0.80	-0.08	-0.08	1.20	0.45	-0.05	-0.01	0.92	0.4808	3.1026	22.1082	1.4083	127
			(-0.62)	(-0.59)	(38.59)	(11.41)	(-1.23)	(-0.21)						
4	-0.2651	0.82	-0.09	-0.04	1.23	0.52	-0.13	-0.08	0.89	0.5124	3.0294	21.8263	1.3622	127
			(-0.54)	(-0.22)	(30.95)	(10.29)	(-2.46)	(-2.62)						
5	-0.0564	1.07	0.16	0.23	1.24	0.65	-0.31	-0.09	0.88	0.5640	3.0346	21.4427	1.3650	127
			(0.83)	(1.14)	(27.22)	(11.10)	(-4.97)	(-2.52)						
5-1	0.7704	0.61	0.49	0.55	0.17	0.37	-0.29	-0.09	0.37	0.1416	-0.6151	-1.2092	-0.2167	
t(5-1)	80.02	2.25	2.23	2.52	3.47	5.72	-4.28	-2.30		38.52	-25.18	-29.03	-13.77	

Table 5.1: Risk-neutral skewness portfolio sorts.

points per month (t-stat: 2.23) and α_{FFC} of 55 basis points per month (t-stat: 2.52). To ensure the return differential is not driven only by stocks in the extreme ends of the RNS cross-sectional distribution, we calculate the performance of the corresponding spread strategy between the highest and the lowest RNS tercile portfolios which is on average 52 basis points per month (t-stat: 2.34), α_{FF} of 42 basis points per month (t-stat: 2.36) and α_{FFC} of 47 basis points per month (t-stat: 2.65). In both cases, the return differential is driven by the underperformance of highly negative RNS portfolio.

We also regress monthly returns on the market factor (MKT), Fama and French (1993) size (SMB) and value (HML) factors and Carhart (1997) momentum (MOM) factor. Table 5.1 reports β_{MKT}, β_{SMB}, β_{HML} and β_{MOM} as well as the explanatory power of the Fama-French-Carhart model. We find that the highest RNS (quintile and tercile) portfolio exhibits significantly higher β_{MKT} and β_{SMB} relative to the lowest RNS portfolio, but it also exhibits significantly lower β_{HML}. Overall, the return differential between the highest and the lowest RNS portfolios cannot be fully explained by these differences in factor loadings, and hence it gives rise to economically and statistically significant alphas.

Finally, for each portfolio we also report the average number of constituent stocks (N) as well as their average value of risk-neutral volatility (RNV), RNS, risk-neutral kurtosis (RNK), average log market value (Ln MV) and average standard deviation of log market value (Std. dev. of ln MV). We find that the lowest RNS portfolios tend to have low RNV and high RNK, whereas portfolios that have the highest RNS tend to have high RNV and low RNK. The results in Table 5.1 also reveal that the lowest (highest) RNS portfolios tend to contain larger (smaller) firms, but there is considerable variation in the MV of firms in each portfolio. Given this set of findings, RNV, RNK and Ln MV are used as control variables in the Fama-MacBeth regressions later.

5.5.2 Robustness Tests

In this subsection, we examine whether our main finding on the relationship between RNS and future stock returns is robust to alternative methodologies. Table 5.2 reports the average portfolio returns, α_{FF} and α_{FFC} for tercile and quintile portfolios as well as the average number of constituent stocks (N) for each of the robustness test.

As a first robustness test, we examine the implications of using an alternative method of estimating RNS. To approximate the integrals (5.3)-(5.5), we follow Conrad, Dittmar and Ghysels (2013) and instead of Simpson's rule, we use the trapezium rule with an equal number of out-of-the-money put and out-of-the-money call options that are closest to the stock price. We also require at least two put and two call options. We find that our main finding is remarkably robust. Panel A of Table 5.2 shows that the portfolio with the highest RNS stocks significantly outperforms the portfolio with the lowest RNS stocks. In the case of quintile portfolios, this differential performance is around 50 basis points per month regardless of the metric used. Again, it is the portfolio with the lowest RNS stocks that yields significantly negative risk-adjusted performance. Therefore, we conclude that the difference between our findings and the results in Conrad, Dittmar and Ghysels (2013) is not due to the method used to calculate RNS.

The second robustness test is to use RNS on the latest available day in the month instead of RNS on the last trading day of the month as a basis for sorting stocks into portfolios. In this way, we include stocks that we may have excluded in the results in the previous section due to missing data on the last trading day of the month. We note that in this approach the number of stocks increases from 635 on average in the benchmark approach to 735 on average. Panel B of Table 5.2 shows that the highest RNS portfolio significantly outperforms the lowest RNS portfolio and this return differential is again driven by the significantly negative performance of the latter portfolio. This differential is highly economically significant for both tercile and quintile portfolios, but the statistical significance is slightly reduced in the case of quintiles.

We then examine whether the relation between RNS and future stock return depends on the weighting scheme used to construct portfolios. Panel C of Table 5.2 shows that the return differential between the highest and the lowest RNS value-weighted portfolios is of the same magnitude as in the benchmark results using equally-weighted portfolios: more than 50 basis points per month in terms of raw returns and around 45 basis points per month in terms of alphas, in all cases considered. This highly economically significant return differential is less statistically significant in the case of quintile portfolios because value-weighted portfolio returns are more volatile relative to the equally-weighted ones.

Finally, in our attempt to explain why our results differ from the ones reported in Conrad, Dittmar and Ghysels (2013), we also sort stocks according to the monthly average of RNS with at least 10 observations per month, instead of using the corresponding estimate on the last trading day of the month. This

Panel A: Alternative method to extract RNS

Terciles	Mean return	α_{FF}	α_{FFC}	N	Quintiles	Mean return	α_{FF}	α_{FFC}	N
1	0.53	-0.34***	-0.30***	212	1	0.41	-0.46***	-0.43***	127
2	0.77	-0.10	-0.06	212	2	0.69	-0.19	-0.13	127
3	0.93	0.05	0.06	212	3	0.90	0.04	0.06	127
3-1	0.40	0.39***	0.36***		4	0.81	-0.08	-0.07	127
t(3-1)	2.84	3.14	2.92		5	0.91	0.05	0.06	127
					5-1	0.50	0.51***	0.49***	
					t(5-1)	2.59	3.02	2.89	

Panel B: Latest available observation of RNS

Terciles	Mean return	α_{FF}	α_{FFC}	N	Quintiles	Mean return	α_{FF}	α_{FFC}	N
1	0.50	-0.31**	-0.30***	245	1	0.51	-0.28**	-0.27**	147
2	0.73	-0.16	-0.15	245	2	0.57	-0.31***	-0.30***	147
3	1.00	0.07	0.15	245	3	0.76	-0.13	-0.11	147
3-1	0.50	0.38**	0.45***		4	0.86	-0.07	-0.01	147
t(3-1)	2.24	2.23	2.66		5	1.02	0.10	0.18	147
					5-1	0.51	0.38*	0.45**	
					t(5-1)	1.92	1.78	2.16	

Panel C: Value-weighted portfolio returns

Terciles	Mean return	α_{FF}	α_{FFC}	N	Quintiles	Mean return	α_{FF}	α_{FFC}	N
1	0.53	-0.09	-0.11	212	1	0.47	-0.14	-0.15	127
2	0.94	0.23*	0.20*	212	2	0.80	0.11	0.08	127
3	1.06	0.33*	0.35*	212	3	0.93	0.23*	0.21	127
3-1	0.53	0.43**	0.46**		4	0.98	0.27	0.28	127
t(3-1)	2.21	1.99	2.12		5	1.06	0.30	0.32	127
					5-1	0.59	0.44	0.47*	
					t(5-1)	2.01	1.61	1.72	

Panel D: Monthly average value of RNS

Terciles	Mean return	α_{FF}	α_{FFC}	N	Quintiles	Mean return	α_{FF}	α_{FFC}	N
1	0.74	-0.06	-0.06	245	1	0.58	-0.21**	-0.19*	147
2	0.70	-0.19	-0.17	245	2	0.91	0.07	0.07	147
3	0.79	-0.15	-0.08	245	3	0.67	-0.23	-0.21	147
3-1	0.04	-0.09	-0.02		4	0.77	-0.14	-0.10	147
t(3-1)	0.17	-0.51	-0.10		5	0.78	-0.17	-0.08	147
					5-1	0.20	0.04	0.11	
					t(5-1)	0.65	0.16	0.52	

Table 5.2: Robustness tests.

is an important modification because daily RNS are highly time-varying. The average AR(1) coefficient of RNS is 0.5958, which indicates that RNS extracted from the daily data is not very persistent. The portfolio performance reported in Panel D of Table 5.2 confirms our conjecture. The return differential between the highest RNS and lowest RNS portfolios disappear, yielding no significant conclusion. This finding highlights the importance of using information on the last trading day of the month, as averaging dilutes the most recent information embedded in the option prices (Rehman and Vilkov (2012)) and explains why our results differ from those in Conrad, Dittmar and Ghysels (2013).

Constructing portfolios on the basis of RNS estimates on the last trading day of each month implies monthly portfolio rebalancing. Therefore, it is worth examining the migration dynamics of stocks across portfolios. In Table 5.3, we analyze the migration of stocks one, two and three months after the initial portfolio formation. We compare the percentage of firms in each quintile in month t assigned to each quintile in month $t+1$, $t+2$ and $t+3$. For example, in Panel A, 50.10% firms in the first RNS tercile in month t were in the same RNS tercile in month $t+1$. If RNS is not available in consecutive months, such stock is assigned to neither category. We find that, on average, stocks do migrate across portfolios but few stocks have their classification switched between the extreme RNS portfolios, as they mainly move to nearby portfolios. Moreover, this migration is not considerably increased as we extend the horizon to two or three months. These results are consistent with the finding that daily RNS estimates are not highly persistent, and hence its most recent value should be used for portfolio formation. However, there is still a considerable degree of persistence in extreme RNS portfolios. The more stable these portfolios, the lesser the need to rebalance the portfolios and the smaller the transaction cost in following the strategy that takes a long position in high RNS portfolio and a short position in low RNS portfolio.

5.5.3 Risk-Neutral Coskewness and Idiosyncratic Skewness Portfolio Sorts

Here, we examine which component of RNS drives the positive relationship between RNS and future realized returns. Using the approach of Conrad, Dittmar and Ghysels (2013) described in Section 5.4, we decompose RNS into its systematic and idiosyncratic components. Then, at the end of each month, firms are

Panel A: Terciles				
(a) Migrations after 1 month				
This/next month	Tercile 1	Tercile 2	Tercile 3	Neither
Tercile 1	50.10%	21.34%	8.66%	19.90%
Tercile 2	21.36%	34.38%	22.83%	21.43%
Tercile 3	8.89%	22.94%	43.27%	24.90%
(b) Migrations after 2 months				
This/after 2 months	Tercile 1	Tercile 2	Tercile 3	Neither
Tercile 1	44.91%	20.76%	9.54%	24.79%
Tercile 2	20.94%	31.04%	21.44%	26.58%
Tercile 3	9.63%	21.80%	37.82%	30.75%
(c) Migrations after 3 months				
This/after 3 months	Tercile 1	Tercile 2	Tercile 3	Neither
Tercile 1	43.58%	20.42%	9.84%	26.17%
Tercile 2	20.56%	30.20%	21.40%	27.84%
Tercile 3	9.78%	21.65%	35.87%	32.69%

Panel B: Quintiles						
(a) Migrations after 1 month						
This/next month	Quntile 1	Quntile 2	Quntile 3	Quntile 4	Quntile 5	Neither
Quntile 1	43.38%	19.20%	8.95%	5.09%	3.65%	19.72%
Quntile 2	18.78%	25.04%	18.95%	10.96%	6.26%	20.00%
Quntile 3	9.38%	18.01%	21.55%	18.40%	10.92%	21.74%
Quntile 4	5.22%	11.60%	18.25%	23.49%	18.39%	23.04%
Quntile 5	3.68%	6.38%	10.83%	19.05%	34.17%	25.88%
(b) Migrations after 2 months						
This/after 2 months	Quntile 1	Quntile 2	Quntile 3	Quntile 4	Quntile 5	Neither
Quntile 1	37.85%	18.09%	9.80%	5.62%	4.12%	24.52%
Quntile 2	17.95%	22.52%	16.65%	11.09%	6.58%	25.22%
Quntile 3	9.74%	16.96%	19.14%	17.08%	10.66%	26.43%
Quntile 4	6.03%	10.99%	16.50%	20.50%	17.25%	28.74%
Quntile 5	4.21%	6.79%	10.98%	17.39%	28.68%	31.96%
(c) Migrations after 3 months						
This/after 3 months	Quntile 1	Quntile 2	Quntile 3	Quntile 4	Quntile 5	Neither
Quntile 1	36.48%	18.04%	9.61%	5.69%	4.24%	25.95%
Quntile 2	17.55%	21.37%	16.72%	10.88%	7.16%	26.32%
Quntile 3	9.74%	16.63%	18.33%	16.64%	10.92%	27.73%
Quntile 4	5.85%	11.33%	16.24%	19.44%	16.50%	30.65%
Quntile 5	4.08%	6.89%	11.27%	17.15%	26.77%	33.84%

Table 5.3: Migration dynamics of portfolios sorted on the basis of risk-neutral skewness.

ranked into terciles (quintiles) based on their RNCS or RNIS in the past month. Each tercile (quintile) contains 33⅓% (20%) of the sample. Finally, we form three (five) equally weighted portfolios on the basis of RNCS or RNIS. Portfolios are formed every month from February 1996 to December 2012.

Table 5.4 reports average values of RNS, average portfolio returns as well as their Fama-French (α_{FF}) and Fama-French-Carhart (α_{FFC}) alphas estimated from the corresponding 3-factor and 4-factor models, during the period 1996-2012. We show in Panel B of Table 5.4 that the tercile portfolio of stocks with the highest RNIS outperforms the portfolio of stocks with the lowest RNIS by 41 basis points per month in terms of raw returns, 32 basis points per month in terms of α_{FF} and 36 basis points per month in terms of α_{FFC}. Similarly, the quintile portfolio of stocks with the highest RNIS outperforms the portfolio of stocks with the lowest RNIS by 62 basis points per month in terms of raw returns, 51 basis points per month in terms of α_{FF} and 57 basis points per month in terms of α_{FFC}. Finally, the return differential is mainly driven by the underperformance of the stocks with the lowest RNIS.

On the other hand, in Panel A of Table 5.4 we find no clear relationship between risk-neutral coskewness (RNCS) and future realized stock returns as the return differential in average raw returns between the extreme terciles or quintiles is insignificant, both economically and statistically. Table 5.4 also reports the average RNS values of stocks across RNCS-sorted and RNIS-sorted portfolios to examine which of the two sorts mimics the steep RNS gradient that we reported across RNS-sorted portfolios in Table 5.1. We find that sorting by RNIS yield a steep gradient of RNS values, while the corresponding gradient for RNCS sorts is much less pronounced. To confirm this finding, we compute the average pairwise correlation between daily RNS and RNIS values across stocks. We find that their correlation is extremely high and equal to 0.896. Overall, these results imply that it is the idiosyncratic component of RNS that drives the positive relationship between RNS and future realized stock returns.

5.5.4 Fama-MacBeth Regressions

To further assess the relationship between RNS and future realized stock returns, we estimate cross-sectional regressions using the approach of Fama and MacBeth (1973). For each month, we estimate cross-sectional regressions of individual stock excess returns on their lagged RNS values controlling for a series of other firm and

Panel A: Risk-neutral coskewness sorts									
Terciles	RNS	Mean return	α_{FF}	α_{FFC}	Quintiles	RNS	Mean return	α_{FF}	α_{FFC}
1	-0.495	0.63	-0.10	-0.07	1	-0.526	0.61	-0.09	-0.07
			(-0.94)	(-0.67)				(-0.78)	(-0.58)
2	-0.401	0.84	-0.05	0.00	2	-0.440	0.73	-0.10	-0.05
			(-0.33)	(-0.03)				(-0.71)	(-0.39)
3	-0.340	0.76	-0.23	-0.22	3	-0.399	0.85	-0.03	0.02
			(-1.55)	(-1.50)				(-0.23)	(0.12)
3-1	0.155	0.13	-0.12	-0.15	4	-0.374	1.04	0.09	0.13
t(3-1)	30.31	0.51	-0.77	-0.91				(0.59)	(0.80)
					5	-0.320	0.50	-0.50	-0.53
								(-2.73)	(-2.85)
					5-1	0.206	-0.11	-0.41	-0.46
					t(5-1)	30.25	-0.33	-1.89	-2.10

Panel B: Risk-neutral idiosyncratic skewness sorts									
Terciles	RNS	Mean return	α_{FF}	α_{FFC}	Quintiles	RNS	Mean return	α_{FF}	α_{FFC}
1	-0.682	0.54	-0.28	-0.27	1	-0.778	0.40	-0.39	-0.40
			(-2.60)	(-2.54)				(-3.11)	(-3.14)
2	-0.388	0.74	-0.13	-0.12	2	-0.513	0.79	-0.07	-0.05
			(-1.08)	(-0.92)				(-0.61)	(-0.43)
3	-0.166	0.96	0.04	0.09	3	-0.388	0.77	-0.09	-0.09
			(0.22)	(0.54)				(-0.67)	(-0.63)
3-1	0.516	0.41	0.32	0.36	4	-0.279	0.74	-0.18	-0.13
t(3-1)	72.31	1.93	1.89	2.19				(-1.14)	(-0.81)
					5	-0.103	1.03	0.12	0.17
								(0.62)	(0.90)
					5-1	0.676	0.62	0.51	0.57
					t(5-1)	69.07	2.40	2.39	2.68

Table 5.4: Risk-neutral coskewness and idiosyncratic skewness portfolio sorts.

stock return characteristics. Table 5.5 reports the average slope coefficients estimated from these monthly cross-sectional regressions as well as their Newey-West adjusted t-statistics. Model (1) uses only stocks' RNS values as an explanatory variable, documenting a significant positive relationship between RNS and future realized stock returns. To assess the magnitude of the RNS coefficient, we note that the spread in average RNS values between the highest and the lowest quintiles in Table 5.1 is 0.7704. Therefore, the reported RNS coefficient implies that the average return differential between the extreme RNS quintile portfolios should be equal to 56 basis points(=0.7704*0.0073) per month, which is in line with the spread return reported in Table 5.1. In model (2) we also account for a set of stock characteristics that are commonly regarded as determinants of stocks returns (market beta, size, book-to-market ratio, momentum and reversal). Interestingly, accounting for these characteristics, the positive relationship between RNS and stock return not only remains remarkably robust, but it also becomes much more significant in statistical terms. In other words, we confirm that RNS does not mimic a beta, size, value, momentum or reversal effect.

We further test the robustness of this relationship by including RNV and RNK in model (3). We find that the magnitude of the average RNS coefficient actually increases and becomes even more statistically significant. This finding shows that the RNS effect is not capturing an RNV or RNK effect. Interestingly, similar to Table 5.1, we also find a statistically significant negative (positive) relationship between RNV (RNK) and future stock returns. Finally, in model (4) we account for a series of other stock characteristics such as idiosyncratic volatility, maximum daily return in the past month, expected idiosyncratic skewness and the probability of jackpot return in the following year that have been found to predict stock returns in recent studies. We find that the magnitude of the RNS coefficient remains very robust and its t-statistic is further increased. Since we are using each of these characteristics in conditional double-sorts in Section 5.6, this is an important result showing that RNS does not mimic any of these effects.

To examine whether the portfolio results based on RNCS and RNIS sorts carry through when we use individual stock returns, we re-run the previous regression models substituting RNS with its systematic and idiosyncratic components. The results are reported under models (5)-(8) in Table 5.5. The main findings from this analysis can be summarized as follows. Firstly, we document a significantly positive relationship between RNIS and future stock returns, confirming the portfolio sort results in Table 5.4. Secondly, the magnitude of this relationship remains

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
RNS	0.0073	0.0072	0.0117	0.0148				
	(2.45)	(3.60)	(4.24)	(4.94)				
RNCS					0.0000	0.0000	0.0002	0.0002
					(0.39)	(0.68)	(3.09)	(3.38)
RNIS					0.0058	0.0057	0.0080	0.0086
					(2.25)	(2.83)	(3.42)	(3.95)
RNV			-0.0337	-0.0283			-0.0359	-0.0307
			(-3.93)	(-3.00)			(-4.13)	(-3.23)
RNK			0.0059	0.0078			0.0047	0.0055
			(2.64)	(2.68)			(2.34)	(2.15)
Beta		-0.0024	0.0018	0.0048		-0.0021	0.0025	0.0054
		(-0.54)	(0.49)	(1.27)		(-0.48)	(0.65)	(1.41)
Ln MV		0.0008	-0.0014	-0.0019		0.0009	-0.0011	-0.0016
		(0.89)	(-1.53)	(-1.53)		(1.07)	(-1.24)	(-1.31)
B/M		0.0021	0.0016	-0.0004		0.0022	0.0017	-0.0004
		(1.72)	(1.55)	(-0.53)		(1.74)	(1.55)	(-0.47)
Momentum		0.0020	0.0031	0.0019		0.0021	0.0032	0.0018
		(0.66)	(1.02)	(0.64)		(0.67)	(1.04)	(0.61)
Reversal		-0.0057	-0.0023	0.0011		-0.0056	-0.0028	0.0003
		(-0.74)	(-0.32)	(0.13)		(-0.74)	(-0.39)	(0.03)
IV				-0.1138				-0.0888
				(-0.82)				(-0.63)
Max				0.0205				0.0183
				(0.82)				(0.73)
EIS				-0.0028				-0.0027
				(-1.34)				(-1.31)
Jackpot				-0.3623				-0.3447
				(-0.62)				(-0.58)
Intercept	0.0082	-0.0103	0.0310	0.0387	0.0075	-0.0133	0.0286	0.0387
	(1.50)	(-0.49)	(1.44)	(1.26)	(1.33)	(-0.65)	(1.36)	(1.25)
T	203	203	203	203	203	203	203	203

Table 5.5: Fama-MacBeth regressions.

remarkably robust and its statistical significance increases when we account for a series of stock characteristics. This is true even when we control for RNV and RNK. Thirdly, we find again that RNV (RNK) is negatively (positively) related to future stock returns. Finally, the coefficient of RNCS is negligible and insignificant in models (5) and (6) but it becomes statistically significant in models (7) and (8), though still of very small magnitude. Taken together, these results show that the positive relationship between RNS and stock returns is driven by its idiosyncratic component, confirming the portfolio sort results in Table 5.4.

5.6 Why Do Highly Negative Risk-Neutral Skewness Stocks Subsequently Underperform?

In this section, we examine the sources of the positive relationship between RNS and future realized stock returns that we documented in Section 5.5. To understand this relationship, we draw insights from the evidence provided in Bollen and Whaley (2004), the sequential trade model of Easley, O'Hara and Srinivas (1998) and the demand-based option pricing model of Garleanu, Pedersen and Poteshman (2009).

In particular, the mechanism we put forward and subsequently test assumes that for some investors certain stocks are relatively overvalued and also too costly or too risky to sell (short-sell). Therefore, investors resort to the options market, buying protective puts on these overvalued stocks as a means of hedging their positive delta exposure and/or trading on their negative expectations. Since risk-averse market makers cannot perfectly hedge their options positions in the stock market due to short-selling constraints, this hedging demand drives up (down) prices for OTM puts (calls), leading to highly negative RNS (Garleanu, Pedersen and Poteshman (2009)). In this way, option prices contain information that it is not embedded in stock prices, consistent with the sequential trade model of Easley, O'Hara and Srinivas (1998). As this information is slowly diffused to the stock market over time, relatively overpriced stocks subsequently yield negative returns, giving rise to a positive relationship between RNS and future realized stock returns.

For this conjectured mechanism to be valid, three conditions are necessary to hold. First, stocks characterized by higher hedging demand should exhibit more negative RNS; otherwise the demand-based arguments of Bollen and Whaley (2004) and Garleanu, Pedersen and Poteshman (2009) do not hold. Second, the underperformance of stocks with the lowest RNS should be driven by stocks that are simultaneously perceived as relatively overpriced; if there is no perceived overpricing, there are no pessimistic expectations to be traded in the options market. Third, the underperformance of stocks with the lowest RNS should be driven by stocks that are too costly or too risky to sell short; otherwise the perceived overpricing would be quickly corrected in the stock market and investors would not resort to the options market. The following subsections test the validity of each of these three conditions.

5.6.1 The Role of Investor Hedging Demand

The first necessary condition for the conjectured mechanism to hold is that the stocks characterized by high hedging demand exhibit relatively lower RNS. We use four proxies for investor hedging demand: the ratio of aggregate put options volume to all options volume (Taylor, Yadav and Zhang (2009)), the aggregate open interest across all options (Hong and Yogo (2012)), and following Acharya, Lochstoer and Ramadorai (2013), the ratio of CEO stock holdings to base salary and the $Z-$score of Zmijewski (1984) that captures default risk. Using each of these hedging demand proxies, we sort stocks in our sample into tercile and quintile portfolios and calculate their median RNS. Table 5.6 shows the time-series average of these median RNS estimates for tercile (Panel A) and quintile portfolios (Panel B) containing the stocks characterized by the highest and the lowest investor hedging demand respectively. We find that for all four proxies, the portfolio of stocks exhibiting the highest hedging demand have significantly lower median RNS in comparison to the portfolio of stocks exhibiting the lowest hedging demand.

5.6.2 The Role of Relative Overvaluation

The second necessary condition for the conjectured mechanism to hold is that the underperformance of the most negative RNS portfolio is driven by stocks that are perceived by some investors as relatively overpriced. We use four proxies for relative overvaluation: market-to-book ratio (Fama and French (1992)), expected idiosyncratic skewness, (Boyer, Mitton and Vorkink (2010)), maximum daily return (Bali, Cakici and Whitelaw (2011)) and the probability of jackpot return in the following year (Conrad, Kapadia and Xing (2014)). High values for each of these proxies capture relative overpricing, because over-optimistic investors or investors with preference for their lottery-like payoffs have temporarily driven up

Panel A: Terciles				
Tercile	Put/All	Open Interest	CEO	Z
Low hedging	-0.3867	-0.3903	-0.4054	-0.3571
High hedging	-0.4048	-0.4105	-0.4478	-0.4029
3-1	-0.0182	-0.0202	-0.0425	-0.0458
t(3-1)	-6.94	-6.72	-21.18	-22.85
Panel B: Quintiles				
Quintile	Put/All	Open Interest	CEO	Z
Low hedging	-0.3933	-0.3840	-0.3895	-0.3544
High hedging	-0.4564	-0.4493	-0.4380	-0.4021
5-1	-0.0630	-0.0653	-0.0485	-0.0477
t(5-1)	-33.21	-8.85	-19.07	-16.33

Table 5.6: Time-series averages of median risk-neutral skewness for portfolios sorted on the basis of hedging demand proxies: put/all option volume ratio (Put/All), open interest, CEO stock holdings to base salary ratio (CEO) and the Zmijewski (1984) score (Z).

	MABA	EIS	Max	Jackpot
MABA	1	0.0282	0.0455	0.0508
EIS	0.0282	1	0.2167	0.4424
Max	0.0455	0.2167	1	0.3288
Jackpot	0.0508	0.4424	0.3288	1

Table 5.7: Time-series averages of correlations between overvaluation proxies: market-to-book ratio (MABA), the expected idiosyncratic skewness (EIS), the maximum daily stock return in the past month (Max) and the probability of jackpot return in the following year (Jackpot).

prices of these stocks, which leads to a subsequent underperformance as the price correction mechanism reverses the overvaluation.

The time-series averages of correlations between overvaluation proxies can be found in Table 5.7, showing that these are only moderately positively correlated, potentially capturing different dimensions of potential overvaluation.

To test our conjecture, we sort portfolios on the basis of RNS and relative overvaluation proxy. First, at the end of each month, firms are ranked into terciles based on RNS in the past month. We then form three equally weighted portfolios on the basis of RNS. After that, we rank firms within each RNS sorted tercile portfolio into terciles based on the relative overvaluation proxy. Finally, we form three equally weighted portfolios on the basis of the firm characteristic within each tercile RNS sorted tercile portfolio. Portfolios are formed every month from February 1996 to December 2012. We repeat this process four times for each pair of the RNS and relative overvaluation proxy to obtain four sets of 9 portfolios.

The results are reported in Table 5.8.

The excess return for low RNS, high market-to-book ratio portfolio is on average -56 basis points per month, whereas the excess return for high RNS, low book-to-market ratio portfolio is on average 13 basis points per month. Similarly, the excess return for low RNS, high expected idiosyncratic skewness portfolio is on average -58 basis points per month, whereas the excess return for high RNS, low expected idiosyncratic skewness portfolio is on average 38 basis points per month. Likewise, the excess return for low RNS, high maximum daily return portfolio is on average -77 basis points per month, whereas the excess return for high RNS, low maximum daily return portfolio is on average 41 basis points per month. Finally, the excess return for low RNS, high probability of achieving a jackpot return stocks portfolio is on average -60 basis points per month, whereas the excess return for high RNS, low probability of achieving a jackpot return stocks portfolio is on average 26 basis points per month. Hence, we confirm that the underperformance of the most negative RNS portfolio is driven by the relatively overpriced stocks.

5.6.3 The Role of Short-Selling Constraints

The third necessary condition for the conjectured mechanism to hold is that the underperformance of the most negative RNS portfolio is driven by stocks that are too costly or too risky to sell short. We use four proxies to capture short-selling constraints: estimated shorting fee (Boehme Danielsen and Sorescu (2006)), relative short interest (Asquith Pathak and Ritter (2005)), idiosyncratic volatility (Wurgler and Zhuravskaya (2002)) and size. High values for the first three proxies and small firm size typically reflect severe short-selling constraints. To test our conjecture, we sort portfolios on the basis of RNS and short-selling constraints proxy. First, at the end of each month, firms are ranked into terciles based on RNS in the past month. We then form three equally weighted portfolios on the basis of RNS. After that, we rank firms within each RNS sorted tercile portfolio into terciles based on the short-selling constraints proxy. Finally, we form three equally weighted portfolios on the basis of the firm characteristic within each RNS sorted tercile portfolio. Portfolios are formed every month from February 1996 to December 2012. We repeat this process four times for each pair of the RNS and short-selling constraints proxy to obtain four sets of 9 portfolios. Since we use conditional sorts, the arguments regarding short-selling

Panel A: MABA				
	MABA 1	MABA 2	MABA 3	Difference
RNS 1	-0.15	-0.37***	-0.56***	-0.40*
	(-0.90)	(-2.62)	(-2.85)	(-1.84)
RNS 2	-0.07	-0.22	-0.16	-0.08
	(-0.35)	(-1.24)	(-0.75)	(-0.31)
RNS 3	0.13	0.24	0.14	0.00
	(0.50)	(1.00)	(0.61)	(0.01)
Panel B: EIS				
	EIS 1	EIS 2	EIS 3	Difference
RNS 1	0.08	-0.30**	-0.58***	-0.67***
	(0.59)	(-2.03)	(-3.19)	(-2.98)
RNS 2	0.21	-0.03	-0.45***	-0.66***
	(1.27)	(-0.20)	(-2.61)	(-2.80)
RNS 3	0.38*	0.14	-0.24	-0.63**
	(1.84)	(0.58)	(-1.09)	(-2.50)
Panel C: Max				
	Max 1	Max 2	Max 3	Difference
RNS 1	0.01	-0.22	-0.77***	-0.78***
	(0.07)	(-1.60)	(-3.36)	(-2.85)
RNS 2	0.38***	-0.12	-0.68***	-1.06***
	(2.84)	(-0.70)	(-3.06)	(-4.06)
RNS 3	0.41**	0.08	-0.21	-0.62**
	(2.27)	(0.38)	(-0.78)	(-2.16)
Panel D: Jackpot				
	Jackpot 1	Jackpot 2	Jackpot 3	Difference
RNS 1	-0.06	-0.23	-0.60***	-0.54**
	(-0.55)	(-1.44)	(-2.77)	(-2.26)
RNS 2	0.25	-0.01	-0.48**	-0.73***
	(1.59)	(-0.08)	(-2.12)	(-2.89)
RNS 3	0.26	0.37	-0.04	-0.30
	(1.15)	(1.47)	(-0.17)	(-1.18)

Table 5.8: Excess returns for double-sorted portfolios on the basis of risk-neutral skewness (RNS) and overvaluation proxies: market-to-book ratio (MABA), the expected idiosyncratic skewness (EIS), the maximum daily stock return in the past month (Max) and the probability of jackpot return in the following year (Jackpot).

constraints and size are relative in nature. We argue that within the lowest RNS portfolio there are smaller firms, and hence more short selling constrained/ higher idiosyncratic risk firms, relative to the other firms in the same RNS portfolio. These firms are not necessarily small in absolute terms. The results of conditional sorts are reported in Table 5.9.

The excess return for low RNS, high estimated shorting fee portfolio is on average -64 basis points per month, whereas the excess return for high RNS, low estimated shorting fee portfolio is on average 43 basis points per month. Similarly, the excess return for low RNS, high relative short interest portfolio is on average -72 basis points per month, whereas the excess return for high RNS, low relative short interest portfolio is on average 41 basis points per month. Likewise, the excess return for low RNS, small stocks portfolio is on average -74 basis points per month, whereas the excess return for high RNS, large stocks portfolio is on average 31 basis points per month. Finally, the excess return for low RNS, high idiosyncratic volatility portfolio is on average -67 basis points per month, whereas the excess return for high RNS, low idiosyncratic volatility portfolio is on average 26 basis points per month. Hence, we confirm that the underperformance of the most negative RNS portfolio is almost exclusively driven by stocks that are too costly or too risky to sell short.

5.7 Conclusion

In this chapter, we use a sample of options on individual stocks and estimate the risk-neutral skewness (RNS) according to the model-free methodology of Bakshi, Kapadia and Madan (2003). We document positive relationship between the RNS and future realized stock returns. A strategy that is long the quintile portfolio with the highest RNS stocks and short the quintile portfolio with the lowest RNS stocks, rebalanced monthly, yields an average return of 61 basis points (t-stat: 2.25) per month, Fama-French (FF) alpha of 49 basis points (t-stat: 2.23) per month and Fama-French-Carhart (FFC) alpha of 55 basis points (t-stat: 2.52) per month. In this case, the return of this strategy is mainly driven by the underperformance of highly negative RNS portfolio. We also confirm positive relation between RNS and future realized stock returns using Fama-MacBeth (1973) regressions after controlling for a series of other firm characteristics that have been shown to predict future stock returns. We further decompose RNS into its systematic and idiosyncratic components and demonstrate that the idiosyncratic

Panel A: ESF				
	ESF 1	ESF 2	ESF 3	Difference
RNS 1	0.03	0.04	-0.64***	-0.67***
	(0.21)	(0.32)	(-3.12)	(-2.99)
RNS 2	0.11	-0.02	-0.47**	-0.59***
	(0.76)	(-0.16)	(-2.25)	(-2.88)
RNS 3	0.43**	0.32	-0.54**	-0.98***
	(2.04)	(1.34)	(-2.27)	(-4.06)
Panel B: RSI				
	RSI 1	RSI 2	RSI 3	Difference
RNS 1	0.21*	-0.07	-0.72***	-0.93***
	(1.78)	(-0.47)	(-3.42)	(-4.17)
RNS 2	0.24	-0.13	-0.50**	-0.74***
	(1.63)	(-0.74)	(-2.42)	(-3.40)
RNS 3	0.41*	0.27	-0.46**	-0.87***
	(1.90)	(1.09)	(-2.00)	(-3.48)
Panel C: Size				
	Large	Medium	Small	Difference
RNS 1	-0.06	-0.19	-0.74***	0.67***
	(-0.71)	(-1.25)	(-3.84)	(3.51)
RNS 2	0.19	0.09	-0.70***	0.89***
	(1.65)	(0.50)	(-3.53)	(4.04)
RNS 3	0.31	0.04	-0.07	0.38
	(1.45)	(0.18)	(-0.31)	(1.47)
Panel D: IV				
	IV 1	IV 2	IV 3	Difference
RNS 1	0.00	-0.32**	-0.67***	-0.67**
	(0.04)	(-2.25)	(-2.93)	(-2.53)
RNS 2	0.37***	-0.13	-0.65***	-1.02***
	(3.18)	(-0.81)	(-2.74)	(-3.67)
RNS 3	0.26	0.25	-0.23	-0.49*
	(1.52)	(1.05)	(-0.89)	(-1.78)

Table 5.9: Excess returns for double-sorted portfolios on the basis of risk-neutral skewness (RNS) and short-sale constraints measures: estimated shorting fee (ESF), relative short interest (RSI), size and idiosyncratic volatility (IV).

RNS is positively related to the future returns. A strategy that is long the quintile portfolio with the highest RNIS stocks and short the quintile portfolio with the lowest RNIS stocks, rebalanced monthly, yields an average return of 62 basis points (t-stat: 2.40) per month, Fama-French (FF) alpha of 51 basis points (t-stat: 2.39) per month and Fama-French-Carhart (FFC) alpha of 57 basis points (t-stat: 2.68) per month.

To explain the positive relationship between RNS and future realized stock returns, we put forward a mechanism on the basis that certain stocks are relatively overpriced and too costly or too risky to sell (short-sell). Hedgers (speculators) cannot sell (short-sell) the stocks and resort to the options market buying OTM puts and/or selling OTM calls. Since risk-averse market makers cannot perfectly hedge their options positions in the stock market due to the stock short-selling constraints, the hedging demand drives up (down) prices for OTM puts (calls), leading to highly negative RNS (Garleanu, Pedersen and Poteshman (2009)). In this way, option prices contain information that it is not embedded in stock prices, consistent with the sequential trade model of Easley, O'Hara and Srinivas (1998). As this information is diffused to the stock market over time, relatively overpriced stocks subsequently yield negative returns, giving rise to a positive relationship between RNS and future realized stock returns.

Our empirical tests confirm the validity of this mechanism. Firstly, stocks characterized by higher hedging demand exhibit, on average, more negative RNS. Secondly, the underperformance of stocks with the lowest RNS is driven by those stocks that are simultaneously characterized as relatively overpriced. Thirdly, the underperformance of stocks with the lowest RNS is driven by those stocks that are also too costly or too risky to sell short. In conclusion, consistent with the sequential trade model of Easley, O'Hara and Srinivas (1998), our findings confirm that option prices contain valuable information that leads stock prices. Therefore, option-implied information should be also considered in asset pricing studies, as in the recent study of An et al. (2014).

Chapter 6

Conclusion

This thesis investigates derivative pricing and risk management under stochastic volatility in a Monte Carlo framework and the relationship between the risk-neutral skewness (RNS) of the returns distribution and future realized stock returns.

The first contribution of the thesis is an effective numerical scheme for importance sampling scheme of Fouque and Tullie (2002), which is computationally expensive as it requires Black-Scholes option price and delta at every time step. Our approach relies on a 2-dimensional lookup table that significantly reduces the number of times Black-Scholes option prices and deltas are calculated. With the lookup table, the effective performance of importance sampling is much better than that of basic Monte Carlo or Control Variate. We have also extended this importance sampling scheme so that it can accommodate more complex and realistic price dynamics such as Heston's stochastic volatility, Bates's stochastic volatility with jumps, Heston-Hull-White and Heston-Cox-Ingersoll-Ross with stochastic volatility and stochastic interest rates. We demonstrate the power of the enhanced importance sampling scheme by pricing single asset options such as European and barrier options as well as multi-asset options such as basket, best-of, worst-of, spread, absolute, composite and quotient options.

Second, we apply importance sampling scheme of Fouque and Tullie (2002) in a Multi-Level Monte Carlo simulation. Here, the improvement comes from lower computational cost of Multi-Level Monte Carlo compared to the basic Monte Carlo and variance reduction through importance sampling.

Third, we show how the Greeks can be computed using the Likelihood Ratio Method based on characteristic function, and how combining it with importance sampling leads to a significant variance reduction for the Greeks.

Finally, we document that RNS is positively related to future realized stock returns. Decomposing RNS into its systematic and idiosyncratic components, we find that the latter drives the positive relationship with future stock returns. To explain the positive relationship between RNS and future realized stock returns, we put forward a mechanism on the basis that certain stocks are relatively overpriced and too costly or too risky to sell (short-sell). As a result, hedgers and speculators resort to the options market as a means of hedging their underlying positions and/or trading on their negative expectations, driving RNS to very low negative values. In this way, option prices contain information that it is not embedded in stock prices. As this information is diffused to the stock market over time, relatively overpriced stocks subsequently yield negative returns, giving rise to a positive relationship between RNS and future realized stock returns. Our empirical tests confirm the validity of this mechanism.

There are several avenues for future research based on the findings of this thesis. First of all, the importance sampling scheme of Fouque and Tullie (2002) could be extended to infinite activity Lévy option pricing models such as that in Barndorff-Nielsen-Shephard (2001). It will be also useful to analyze the finite difference approximation error in estimating the Greeks using the Likelihood Ratio Method based on characteristic function. Another direction for future research is to find the explanation for the negative (positive) relationship between risk-neutral volatility (risk-neutral kurtosis) and future realized stock returns that we find using Fama-MacBeth (1973) regressions.

Bibliography

Acharya, V., Lochstoer, L., and Ramadorai, T. (2013). Limits to arbitrage and hedging: Evidence from commodity markets. *Journal of Financial Economics*, 109(2):441–465.

Amemiya, T. (1985). *Advanced Econometrics*. Harvard University Press, Cambridge, MA.

An, B.-J., Ang, A., Bali, T., and Cakici, N. (2014). The joint cross-section of stocks and options. *Journal of Finance*. Forthcoming.

Arouna, B. (2004). Robbins-monro algorithms and variance reduction in finance. *Journal of Computational Finance*, 7(2):35–62.

Asquith, P., Pathak, P., and Ritter, J. (2005). Short interest, institutional ownership, and stock returns. *Journal of Financial Economics*, 78(2):243–276.

Avramidis, A. (2002). Importance sampling for multimodal functions and applications to pricing exotic options. In Yücesan, E., Chen, C.-H., Snowdon, J. L., and Charnes, J. M., editors, *Proceedings of the 2002 Winter Simulation Conference*, pages 1493–1501. IEEE.

Bakshi, G., Cao, C., and Chen, Z. (1997). Empirical performance of alternative option pricing models. *Journal of Finance*, 52(5):2003–2049.

Bakshi, G., Kapadia, N., and Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1):101–143.

Bakshi, G. and Madan, D. (2000). Spanning and derivative-security valuation. *Journal of Financial Economics*, 55(2):205–238.

Bali, T., Cakici, N., and Whitelaw, R. (2011). Maxing out: Stocks as lotteries and the cross-section of future returns. *Journal of Financial Economics*, 99(2):427–446.

Bali, T. and Hovakimian, A. (2009). Volatility spreads and expected stock returns. *Management Science*, 55(11):1797–1812.

Bali, T., Hu, J., and Murray, S. (2013). Option implied volatility, skewness, and kurtosis and the cross-section of expected stock returns. Working Paper, Georgetown University, Singapore Management University, and University of Nebraska - Lincoln.

Bali, T. and Murray, S. (2013). Does risk-neutral skewness predict the cross-section of equity option portfolio returns? *Journal of Financial and Quantitative Analysis*, 48(4):1145–1171.

Barberis, N. and Huang, M. (2008). Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review*, 98(5):2066–2100.

Barndorff-Nielsen, O. and Shephard, N. (2001). Non-gaussian ornstein-uhlenbeck-based models and some of their uses in financial economics. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(2):167–241.

Barraquand, J. (1995). Numerical valuation of high dimensional multivariate european securities. *Management Science*, 41(12):1882–1891.

Bates, D. (1991). The crash of '87: Was it expected? the evidence from options markets. *Journal of Finance*, 46(3):1009–1044.

Bates, D. (1996). Jumps and stochastic volatility: The exchange rate processes implicit in deutschemark options. *Review of Financial Studies*, 9(1):69–107.

Benhamou, E. (2002). Smart monte carlo: Various tricks using malliavin calculus. *Quantitative Finance*, 2(5):329–336.

Binder, A. and Aichinger, M. (2013). *A Workout in Computational Finance.* Wiley, Chichester.

Boehme, R., Danielsen, B., and Sorescu, S. (2006). Short sale constraints, differences of opinion, and overvaluation. *Journal of Financial and Quantitative Analysis*, 41(2):455–488.

Bollen, N. and Whaley, R. (2004). Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance*, 59(2):711–753.

Bouzoubaa, M. and Osseiran, A. (2010). *Exotic Options and Hybrids: A Guide to Structuring, Pricing and Trading*. Wiley, Chichester.

Boyer, B., Mitton, T., and Vorkink, K. (2010). Expected idiosyncratic skewness. *Review of Financial Studies*, 23(1):169–202.

Broadie, M. and Kaya, Ö. (2004). Exact simulation of option greeks under stochastic volatility and jump diffusion models. In Ingalls, R. G., Rossetti, M. D., Smith, J. S., and Peters, B. A., editors, *Proceedings of the 2004 Winter Simulation Conference*, pages 1607–1615. IEEE.

Brunnermeier, M., Gollier, C., and Parker, J. (2007). Optimal beliefs, asset prices, and the preference for skewed returns. *American Economic Review*, 97(2):159–165.

Capriotti, L. (2007). Least squares importance sampling for libor market models. *Wilmott Magazine,* November, pages 2–11.

Capriotti, L. (2008). Least squares importance sampling for monte carlo security pricing. *Quantitative Finance*, 8(5):485–497.

Chan, J. and Kroese, D. (2010). Efficient estimation of large portfolio loss probabilities in t-copula models. *European Journal of Operational Research*, 205(2):361–367.

Chan, N. and Wong, H.-Y. (2006). *Simulation Techniques in Financial Risk Management*. Wiley-Interscience, Hoboken, NJ.

Chang, B.-Y., Christoffersen, P., Jacobs, K., and Vainberg, G. (2012). Option-implied measures of equity risk. *Review of Finance*, 16(2):385–428.

Chen, Z. and Glasserman, P. (2008). Fast pricing of basket default swaps. *Operations Research*, 56(2):286–303.

Conrad, J., Dittmar, R., and Ghysels, E. (2013). Ex ante skewness and expected stock returns. *Journal of Finance*, 68(1):85–124.

Conrad, J., Kapadia, N., and Xing, Y. (2014). Death and jackpot: Why do individual investors hold overpriced stocks? *Journal of Financial Economics.* Forthcoming.

Cremers, M. and Weinbaum, D. (2010). Deviations from put-call parity and stock return predictability. *Journal of Financial and Quantitative Analysis*, 45(2):335–367.

Dagpunar, J. (2007). *Simulation and Monte Carlo: With applications in finance and MCMC*. Wiley, Chichester.

DeMiguel, V., Plyakha, Y., Uppal, R., and Vilkov, G. (2014). Improving portfolio selection using option-implied volatility and skewness. *Journal of Financial and Quantitative Analysis*. Forthcoming.

Dennis, P. and Mayhew, S. (2002). Risk-neutral skewness: Evidence from stock options. *Journal of Financial and Quantitative Analysis*, 37(3):471–493.

Duan, J.-C. and Wei, J. (2009). Systematic risk and the price structure of individual equity options. *Review of Financial Studies*, 22(5):1981–2006.

Easley, D., O'Hara, M., and Srinivas, P. (1998). Option volume and stock prices: Evidence on where informed traders trade. *Journal of Finance*, 53(2):431–465.

Epps, T. (2007). *Pricing Derivative Securities*. World Scientific, River Edge, NJ.

Étoré, P. and Jourdain, B. (2010). Adaptive optimal allocation in stratified sampling methods. *Methodology and Computing in Applied Probability*, 12(3):335–360.

Fama, E. and French, K. (1992). The cross-section of expected stock returns. *Journal of Finance*, 47(2):427–465.

Fama, E. and French, K. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.

Fama, E. and French, K. (1997). Industry costs of equity. *Journal of Financial Economics*, 43(2):153–193.

Fama, E. and MacBeth, J. (1973). Risk return and equilibrium: Empirical tests. *Journal of Financial Political Economy*, 71:607–636.

Fouque, J. and Han, C. (2004). Variance reduction for monte carlo methods to evaluate option prices under multi-factor stochastic volatility models. *Quantitative Finance*, 4(5):597–606.

Fouque, J., Papanicolaou, G., and Sircar, K. (2000). *Derivatives in Financial Markets with Stochastic Volatility*. Cambridge University Press, Cambridge.

Fouque, J., Papanicolaou, G., and Sircar, K. (2011). *Multiscale Stochastic Volatility for Equity, Interest Rate, and Credit Derivatives*. Cambridge University Press, Cambridge.

Fouque, J. and Tullie, T. (2002). Variance reduction for monte carlo simulation in a stochastic volatility environment. *Quantitative Finance*, 2(1):24–30.

Fournie, E., Lebuchoux, J., and Touzi, N. (1997). Small noise expansion and importance sampling. *Asymptotic Analysis*, 14(4):361–376.

Friesen, G., Zhang, Y., and Zorn, T. (2012). Heterogeneous beliefs and risk neutral skewness. *Journal of Financial and Quantitative Analysis*, 47(4):851–872.

Gao, X. and Ritter, J. (2010). The marketing of seasoned equity offerings. *Journal of Financial Economics*, 97(1):33–52.

Garleanu, N., Pedersen, L., and Poteshman, A. (2009). Demand-based option pricing. *Review of Financial Studies*, 22(10):4259–4299.

Giles, M. (2008). Multi-level monte carlo path simulation. *Operations Research*, 56(3):607–617.

Giles, M. (2009). Multilevel monte carlo for basket options. In Rossetti, M. D., Hill, R. R., Johansson, B., Dunkin, A., and Ingalls, R. G., editors, *Proceedings of the 2009 Winter Simulation Conference*, pages 1283–1290. IEEE.

Glasserman, P. (2003). *Monte Carlo Methods in Financial Engineering*. Springer, New York, NY.

Glasserman, P. (2005). Measuring marginal risk contributions in credit portfolios. *Journal of Computational Finance*, 9(2):1–41.

Glasserman, P., Heidelberger, P., and Shahabuddin, P. (1999a). Asymptotically optimal importance sampling and stratification for pricing path-dependent options. *Mathematical Finance*, 9(2):117–152.

Glasserman, P., Heidelberger, P., and Shahabuddin, P. (1999b). Importance sampling in the heath-jarrow-morton framework. *Journal of Derivatives*, 7(1):32–50.

Glasserman, P., Kang, W., and Shahabuddin, P. (2008). Fast simulation of multifactor portfolio credit risk. *Operations Research*, 56(5):1200–1217.

Glasserman, P. and Li, J. (2005). Importance sampling for portfolio credit risk. *Management Science*, 51(11):1643–1656.

Glasserman, P. and Liu, Z. (2010). Sensitivity estimates from characteristic functions. *Operations Research*, 58(6):1611–1623.

Glasserman, P. and Staum, J. (2001). Conditioning on one-step survival for barrier option simulations. *Operations Research*, 49(6):923–937.

Grundke, P. (2009). Importance sampling for integrated market and credit portfolio models. *European Journal of Operational Research*, 194(1):206–226.

Grzelak, L. and Oosterlee, C. (2011). On the heston model with stochastic interest rates. *SIAM Journal on Financial Mathematics*, 2(1):255–286.

Guasoni, P. and Robertson, S. (2008). Optimal importance sampling with explicit formulas in continuous time. *Finance and Stochastics*, 12(1):1–19.

Guiso, L. and Paiella, M. (2008). Risk aversion, wealth and background risk. *Journal of the European Economic Association*, 6(6):1109–1150.

Han, B. (2008). Investor sentiment and option prices. *Review of Financial Studies*, 21(1):387–414.

Hansis, A., Schlag, C., and Vilkov, G. (2010). The dynamics of risk-neutral implied moments: Evidence from individual options. Working Paper, Goethe University Frankfurt.

Harvey, C. and Siddique, A. (2000). Conditional skewness in asset pricing tests. *Journal of Finance*, 55(3):1263–1295.

Haug, E. (2006). *The Complete Guide to Option Pricing Formulas.* McGraw-Hill, New York, NY, 2 edition.

Heston, H. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343.

Hirsa, A. (2012). *Computational Methods in Finance.* CRC Press, Boca Raton, FL.

Hong, H. and Yogo, M. (2012). What does futures market interest tell us about the macroeconomy and asset prices? *Journal of Financial Economics*, 105(3):473–490.

Huang, P., Subramanian, D., and Xu, J. (2010). An importance sampling method for portfolio cvar estimation with gaussian copula models. In Johansson, B., Jain, S., Montoya-Torres, J., Hugan, J., and Yücesan, E., editors, *Proceedings of the 2010 Winter Simulation Conference*, pages 2790–2800. IEEE.

Jäckel, P. (2002). *Monte Carlo Methods in Finance*. Wiley, Chichester.

Joshi, M. (2005). Applying importance sampling to pricing single tranches of cdos in a one-factor li model. *Wilmott Magazine,* March, pages 82–87.

Joshi, M. and Kainth, D. (2004). Rapid computation of prices and greeks for nth to default swaps in the li model. *Quantitative Finance*, 4(3):266–275.

Joshi, M. and Leung, T. (2007). Using monte carlo simulation and importance sampling to rapidly obtain jump-diffusion prices of continuous barrier options. *Journal of Computational Finance*, 10(4):93–105.

Kang, W. and Shahabuddin, P. (2005). Fast simulation for multifactor portfolio credit risk in the t-copula model. In Kuhl, M. E., Steiger, N. M., Armstrong, F. B., and Joines, J. A., editors, *Proceedings of the 2005 Winter Simulation Conference*, pages 1859–1868. IEEE.

Korn, R., Korn, E., and Kroisandt, G. (2010). *Monte Carlo Methods and Models in Finance and Insurance*. CRC Press, Boca Raton, FL.

Kroese, D., Taimre, T., and Botev, Z. (2011). *Handbook of Monte Carlo Methods*. Wiley, Hoboken, NJ.

Lemieux, C. (2009). *Monte Carlo and Quasi-Monte Carlo Sampling*. Springer, New York, NY.

Merino, S. and Nyfeler, M. (2004). Applying importance sampling for estimating coherent credit risk contributions. *Quantitative Finance*, 4(2):199–207.

Mitton, T. and Vorkink, K. (2007). Equilibrium underdiversification and the preference for skewness. *Review of Financial Studies*, 20(4):1255–1288.

Neddermeyer, J. (2011). Non-parametric partial importance sampling for financial derivative pricing. *Quantitative Finance*, 11(8):1193–1206.

Pellizzari, P. (1998). Efficient monte carlo pricing of basket options. Working Paper, University of Venice.

Rehman, Z. and Vilkov, G. (2012). Risk-neutral skewness: Return predictability and its sources. Working Paper, BlackRock, and Goethe University Frankfurt.

Robertson, S. (2010). Sample path large deviations and optimal importance sampling for stochastic volatility models. *Stochastic Processes and their Applications*, 120(1):66–83.

Ross, S. (2013). *Simulation*. Academic Press, Burlington, MA, 5 edition.

Sak, H. and Hörmann, W. (2012). Fast simulations in credit risk. *Quantitative Finance*, 12(10):1557–1569.

Samuelson, P. (1970). The fundamental approximation theorem of portfolio analysis in terms of means, variances, and higher moments. *Review of Economic Studies*, 37(4):537–542.

Scott, R. and Horvath, P. (1980). On the direction of preference for moments of higher order than the variance. *Journal of Finance*, 35(4):915–919.

Su, Y. and Fu, M. (2002). Optimal importance sampling in securities pricing. *Journal of Computational Finance*, 5(4):27–50.

Tavella, D. (2002). *Quantitative Methods in Derivatives Pricing: An Introduction to Computational Finance*. Wiley, Hoboken, NJ.

Taylor, S., Yadav, P., and Zhang, Y. (2009). Cross-sectional analysis of the risk-neutral skewness. *Journal of Derivatives*, 16(4):38–52.

Toft, K. and Prucyk, B. (1997). Options on leveraged equity: Theory and empirical tests. *Journal of Finance*, 52(3):1151–1180.

Tullie, T. (2002). *Variance Reduction for Monte Carlo Simulation of European, American or Barrier Options in a Stochastic Volatility Environment*. PhD thesis, North Carolina State University.

Vázquez-Abad, F. and Dufresne, D. (1998). Accelerated simulation for pricing asian options. In Medeiros, D. J., Watson, E. F., Carson, J. S., and Manivannan, M. S., editors, *Proceedings of the 1998 Winter Simulation Conference*, pages 1493–1500. IEEE.

Wang, H. (2012). *Monte Carlo Simulation with Applications to Finance*. CRC Press, Boca Raton, FL.

Webber, N. (2011). *Implementing Models of Financial Derivatives: Object Oriented Applications with VBA*. Wiley, Chichester.

Wurgler, J. and Zhuravskaya, E. (2002). Does arbitrage flatten demand curves for stocks? *Journal of Business*, 75(4):583–608.

Xing, Y., Zhang, X., and Zhao, R. (2010). What does individual option volatility smirk tell us about the future equity returns? *Journal of Financial and Quantitative Analysis*, 45(3):641–662.

Zhang, P. (1998). *Exotic Options: A Guide to Second Generation Options*. World Scientific, River Edge, NJ.

Zmijewski, M. (1984). Methodological issues related to the estimation of financial distress prediction models. *Journal of Accounting Research*, 22:59–82.