Derivatives in Dynamic Markets

ISBN 978 90 361 0103 5

Cover design: Crasborn Graphic Designers bno, Valkenburg a.d. Geul

This book is no. 445 of the Tinbergen Institute Research Series, established through cooperation between Thela Thesis and the Tinbergen Institute. A list of books which already appeared in the series can be found in the back.

Derivatives in Dynamic Markets

Derivaten in dynamische markten

Thesis

to obtain the degree of Doctor from the
Erasmus University Rotterdam
by command of the rector magnificus
Prof.dr. S.W.J. Lamberts
and in accordance with the decision of the Doctorate Board

The public defence shall be held on
Thursday 29th January 2009 at 11.00 hrs
by

Amy Sau Kwun Wong
born in Groningen

ERASMUS UNIVERSITEIT ROTTERDAM

Doctoral Committee

Promoters:	Prof.dr. C.G. de Vries
	Prof.dr. D.J.C. van Dijk
Other members:	Prof.dr. P.J.F. Groenen
	Dr. R. Huisman
	Prof.dr. W.F.C. Verschoor

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Financial markets can be very turbulent. Future market developments are hard to predict and there is only one thing certain: The markets will change over time. If the changes in dynamic markets are unexpected by investors, they can encounter substantial risk on their investments. Risk on investments can arise from different sources and recent market developments have shown that correlation risk and liquidity risk can have a considerable impact on financial assets. This thesis investigates the implications of time-varying market conditions for financial derivatives regarding correlation risk and liquidity risk. In practice, many investments involve multiple assets and the correlations between these assets are very important for pricing, hedging and risk management of financial derivatives. Liquidity risk of an investment arises if the asset cannot be traded against its model value at a certain time due to, e.g., unfavorable market conditions.

Investors can diminish their investment risk by diversifying over multiple financial assets. A well-diversified asset portfolio is relatively less risky than investing in only one asset, since a loss in one asset can be offset by a gain in another asset in the portfolio. In order to achieve diversification it is important to choose assets that have low correlations in the portfolio. Correlations are highly time varying (see, e.g., Goetzmann, Li and Rouwenhorst 2005) and changes in the correlations between the underlying assets can have a significant impact on the risk of a financial derivative. A common financial derivative with multiple underlying assets is the basket option. The basket option is an option on a basket of underlying assets and its value is highly dependent on the correlations between the underlying assets in the basket. Chapter 2 investigates the impact of correlation term structures for different market conditions on the basket option and the relevant risk measures.

Determining the correlation values between financial assets is a challenge, since correlations are unobservable in the market and have to be estimated by models. Therefore, correlation values are subject to estimation errors. A stylized fact is that diversification is least available when most needed, since correlations often increase during financial crises. Therefore, one of the most practically relevant correlation term structures is one containing a sudden jump upward. In order to solely measure the impact of different market conditions, the correlation term structures in Chapter 2 do not contain correlation errors. However, different correlation models can render different correlation estimates. Therefore, Chapter 3 combines the above mentioned issues and investigates the size and sign of the risk measure misestimation given a certain size of correlation error for the correlation term structure containing a sudden jump upward.

Chapter 4 discusses the liquidity risks for static hedges of barrier options. Barrier options are quite popular, since they are cheaper than the standard European options and still give the same option payoff for certain ranges of the underlying asset value. Barrier options have an option life that depends on the value of the underlying asset, i.e. these are path-dependent options. For example, if the underlying asset hits the specified barrier at a certain time before the option maturity, the down-and-out barrier option becomes worthless; in all other cases the barrier option has the payoff of a standard European option. A static hedge of barrier options provides a better hedge than a dynamic delta hedge, since higher order option sensitivities are hedged as well by the static hedge (see, e.g., Nalholm and Poulsen (2006a)). The liquidity risk arises when the barrier option knocks out before maturity and the conditions at the time one has to unwind the static hedge are not favourable. Chapter 4 investigates different liquidity conditions and the corresponding liquidity costs for the static hedge.

Chapters 2 and 3 discuss that accurate correlation values and errors have a large impact in risk management of basket options. Correlation estimates highly depend on the model used. For practical applications correlation models need to be parsimonious, such that the number of parameters to be estimated remains manageable for large asset portfolios. A model also needs to be flexible enough to be able to incorporate recent changes in market conditions. The performance of a relatively new and simple correlation model has been tested in Chapter 5 for different correlation scenarios as well as time-varying volatilities.

In the literature, market parameters are often taken to be constant in models for simplicity. However, the reality shows that we live in turbulent times where financial market conditions are not constant — they are constantly changing. The chapters in this thesis address topics on time-varying financial parameters and the results show that parameter time variation is an important aspect that should be taken into account for many financial applications.

Acknowledgements

I would like to thank Casper de Vries and Dick van Dijk for being my promoters and the committee for reading this thesis. Financial support from the Tinbergen Institute and the Econometric Institute is gratefully acknowledged.

Moreover, I thank the fellow PhD students who have contributed to a very pleasant environment. I really enjoyed the social activities together, where we never seemed to have a shortage of discussion topics on and beyond science. Finally, I am greatly indebted to the people who have stood by me through all seasons with their unrelenting support and patience. I will always be grateful for having you in my life.

Amy Wong
Rotterdam, August 2008

Chapter 2

Basel II and Basket Options with Time-varying Correlations[1]

2.1 Introduction

An important aspect of the "International Convergence of Capital Measurement and Capital Standards: A Revised Framework," also known as Basel II, is the management of the risk of financial products that cannot be entirely captured by value-at-risk (VaR); for example, nonlinear products or sudden correlation shifts (see Bank for International Settlements 2005). Basket options are derivatives that belong to the class of products that are subject to nonlinear and correlation risk. A basket option is an option on a portfolio of underlying assets, and the option price is highly dependent on the correlations between the underlying assets. This study examines basket options, because they are widely used across many financial markets, such as foreign exchange markets (see, e.g., Bennett and Kennedy 2004, Dammers and McCauley 2006), credit derivatives markets (e.g., Duffie and Singleton 2003), and equity markets (e.g., Pellizzari 2005). Many studies, such as Margrabe (1978), Curran (1994), Milevsky and Posner (1998), Brigo et al. (2004), and Deelstra, Liinev, and Vanmaele (2004), value basket options under the assumption of constant correlations between the processes of the underlying assets. However, recent empirical studies (e.g., Goetzmann, Li, and

[1]This chapter is a slightly modified version of Wong (2006), that has been published in the International Journal of Central Banking.

Rouwenhorst 2005), have shown that correlations of stock returns are considerably time varying.

Therefore, this chapter examines the impact of time-varying correlation term structures on pricing and hedging of basket options as well as the implications for risk management. Empirical examination of the correlations between equity indices S&P 500, FTSE 100, and the Merrill Lynch government bond index shows that empirical features such as jumps, regime switches, and (nearly) linearly changing correlations can occur in practice. The main contribution of this study is to take these features into account in the correlation term structures of the basket options. To my knowledge, the impact of correlation jumps, regime switches, and linear correlation changes on the value-at-risk of basket options and the performance of the square-root-of-time rule have not yet been examined in the literature. Studies such as Skintzi, Skiadopoulos, and Refenes (2005) have analyzed the impact of estimation errors of constant correlation on value-at-risk of a portfolio of standard European options. Pellizzari (2005) examines a linearly increasing volatility structure for hedged basket options and the corresponding risk measures but does not look at time variation of correlations. Kupiec (1998) performs stress tests taking time-varying correlations into account for portfolios with linear exposure to underlying assets.

A second contribution of this study is an analysis on VaR of hedged basket options and the performance of the square-root-oftime rule for these positions when the correlation term structure contains the above-mentioned time variation. In practice, financial institutions often hedge their outstanding option positions to reduce the exposure to risk of the position and apply the square-root-oftime rule to the one-day VaR in order to obtain an estimate of the regulatory ten-day VaR. The performance of the square-root-oftime rule has been studied, for example, for GARCH processes by Diebold et al. (1997) and for jump diffusion processes by Danielsson and Zigrand (2005).

Moreover, this chapter discusses the differences between the risk measure VaR and the coherent risk measure CVaR (conditional value at risk) for basket options, where the latter (more robust) measure can give additional information needed in some cases to give an adequate risk assessment of the derivatives position. In this chapter, a Monte Carlo simulation study has been performed for basket options with time-varying correlation term structures against the benchmark of constant correlations. The results are as follows.

First, there is an asymmetric correlation effect on the VaR of the basket option, where a change in negative (constant) correlations between the underlying assets has a greater impact on the VaR than a change in positive correlations of the same magnitude. This result is surprising: it is widely known that well-diversified asset portfolios (i.e., negative correlations) are less risky than portfolios with highly correlated assets, so one might expect that a basket option on a well-diversified portfolio is also relatively less risky. The results show that the potential loss *values* given by (C)VaR are indeed lower for basket options on negatively correlated assets, but at the same time the *changes* in (C)VaR are more subject to correlation risk as well. Ignoring this result can lead to serious underestimation of the VaR if sudden changes in market conditions occur. Another implication of this result is that VaR estimation of basket options on well-diversified baskets is relatively more prone to model risk, since in practice correlations have to be estimated under the assumption of a certain correlation model such as RiskMetrics^{TM}. For studies on the impact of estimation errors of constant correlations, see Fengler and Schwendner (2004) for basket options and Skintzi, Skiadopoulos, and Refenes (2005).

Second, the time at which correlation shocks occur during the life of the option is important for the VaR of basket options (especially if hedging is applied), even though the payoff of the basket option only depends on the value of the underlying assets at maturity. Compared with constant correlations at the average value of the time-varying correlations over the life of the option as a benchmark, the results show that the VaR estimates are dependent on the specific type of correlation time variation, even if the average correlation over the life of the option is the same. The estimates for the risk measures are also highly dependent on the hedge effectiveness for the option. This result is relevant for financial institutions, because option positions are often hedged in practice.

Third, the VaR estimate obtained by the square-root-of-time rule can lead to underestimation of the ten-day VaR for the unhedged option with time-varying correlations. The risk assessment of deep out-of-the-money derivatives plays an important role in the Basel II framework. This study shows that the square-root-of-time rule for the risk measure VaR underestimates the risks when the OTM basket option has a (time-varying) highly negatively correlated portfolio, and this underestimation even exceeds the minimum regulatory stress factor of value 3 for some cases. There is also underestimation of the VaR for basket options with constant correlations, but this underestimation remains below a factor of 3. When the time-varying correlations are relatively low at the start of the option, by using the square-root-of-time rule one implicitly assumes that the correlations remain this low for ten days. As a result, the risk implied by the ten-day VaR can be much larger than the estimate obtained from the square-root-of-time rule for time-varying correlations. For hedged options, the performance of the square-root-of-time rule is highly dependent on the difference in hedge effectiveness over a one-day and ten-day horizon and can lead to large deviations from the ten-day VaR.

Finally, VaR gives information about the potential loss of the option position corresponding to a certain confidence level but does not reveal the size of the loss if the VaR is exceeded. The difference in VaR and CVaR can be more than 40 percent for certain correlation term structures. Therefore, it is advisable to use VaR together with the coherent risk measure CVaR, because the CVaR can provide additional information needed to assess the risk of the basket options.

This chapter is structured as follows. Section 2.2 contains an empirical examination of the correlations between equity and bond indices. Section 2.3 proceeds with the simulation framework, and in Section 2.4 the results of the simulation study are discussed. Finally, Section 2.5 concludes.

Figure 2.1: Correlations generated by RiskMetrics^{TM} model

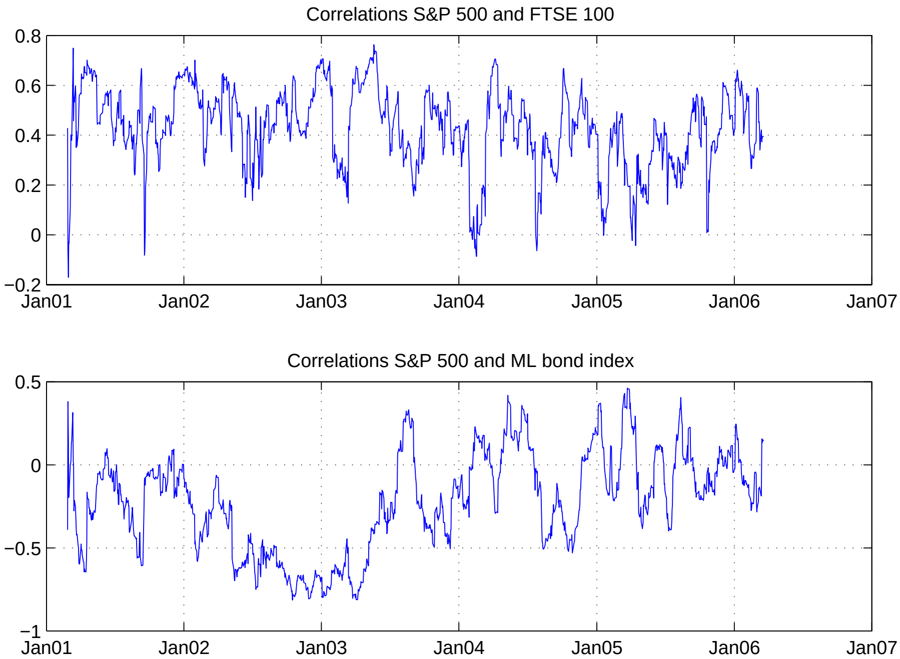

2.2 Empirical motivation

This study examines the impact of time-varying correlations on the basket option and the implications for risk management. In order to see in what way this time variation could occur, the correlations between equity indexes S&P 500 (United States) and FTSE 100 (United Kingdom) as well as the Merrill Lynch US Treasury one- to ten-years bond index (henceforth, ML bond index) are illustrated in Figure 2.1. The data are from October 2000 to March 2006 and are obtained from Datastream. The correlations between the returns of N assets can be estimated in many different ways, but a widely used benchmark model is RiskMetrics^{TM} of J.P. Morgan (1996) specified as follows. Let r_t be the $1 \times N$ vector of asset returns at time t and $H_t = \{\sigma_{ij,t}\}_{i,j=1}^{N}$ be the conditional covariance matrix at time t. Then, RiskMetrics^{TM} estimates the covariance matrix by

$$H_t = \lambda H_{t-1} + (1-\lambda) r'_{t-1} r_{t-1},$$

where the weighting parameter λ has value 0.94 for daily data[2]. Subsequently, the correlations can be easily calculated as $\rho_{ij,t} = \frac{\sigma_{ij,t}}{\sqrt{\sigma_{ii,t}\sigma_{jj,t}}}$.

Figure 2.1 shows that the correlations between S&P 500 and FTSE 100 as well as the correlations between S&P 500 and the ML bond index, are considerably time varying. The correlations between S&P 500 and FTSE 100 remain positive during most of the sample period, but there are several large jumps downward — for example in 2005. The correlations between S&P 500 and the ML bond index are, during some periods, positive (value stays below 0.5) but remain negative most of the time. One can also observe the following correlation patterns in Figure 2.1: starting from the second quarter of 2003 the correlation appears to increase almost linearly. Moreover, one can see from the figure that the correlations from the second quarter of 2002 until the first quarter of 2003 are on average much more negative than the correlations before and after this period, which resembles a regime switch. Following these observations, the correlation term structures in the subsequent simulation study will contain jumps and regime switches as well as linear increase and decrease in correlations during the life of the option.

2.3 Simulation Methods

For simplicity, we will use a basket option on two underlying assets to analyze the impact of the correlation changes on basket options. The asset price dynamics will be simulated using the differential equations

$$\begin{aligned} dS_{1,t} &= \mu_1 S_{1,t} dt + \sigma_1 S_{1,t} dW_{1,t}, \\ dS_{2,t} &= \mu_2 S_{2,t} dt + \sigma_2 S_{2,t} dW_{2,t}, \\ dW_{1,t} dW_{2,t} &= \rho_t dt, \qquad \text{for } t = 1, \ldots, T \quad , \end{aligned} \tag{2.1}$$

where $W_{1,t}$ and $W_{2,t}$ are correlated Brownian motions with correlation ρ_t at time t.

[2]The first 100 daily observations of the data sample are used to estimate the initial covariance matrix.

The value V_t of the basket option with strike price K at time t is given by

$$C_T = (\omega_1 S_{1,T} + \omega_2 S_{2,T} - K)^+ = \max(\omega_1 S_{1,T} + \omega_2 S_{2,T} - K, 0), \quad (2.2)$$

$$V_t = e^{-r(T-t)} \mathbb{E}_Q \left[\sum_{m=1}^{M} C_T^{(m)} / M \right] \quad (2.3)$$

where C_T is the claim of the option at maturity T, ω_i are the portfolio weights of the underlying assets ($i = 1, 2$), Q is the risk-neutral martingale measure, and M is the number of simulations. The delta, Δ_i, is the sensitivity of the option value with respect to the price of asset i (used for delta hedging), and it is computed by the central difference method (for more details, see, e.g., Glasserman 2004):

$$\Delta_i(S_i) = \frac{\partial V}{\partial S_i} \approx \frac{V(S_i + h) - V(S_i - h)}{2h}, \quad i = 1, 2. \quad (2.4)$$

Table 2.1: Option parameters

VaR horizon	10/252	$[\mu_1, \mu_2]$	[0.1,0.1]
Maturity option T	63/252	r	0.05
dt	1/252	K_{ITM}	95
Nr. of assets	2	K_{ATM}	100
$[S_{1,0}, S_{2,0}]$	[100,100]	K_{OTM}	105
$[\sigma_1, \sigma_2]$	[0.35,0.35]	h in (2.4)	0.01
$[\omega_1, \omega_2]$	[1/2,1/2]	M	5000

We will analyze the different correlation specifications for in-the-money (ITM), at-the-money (ATM), and out-of-the-money (OTM) basket options. Delta hedging will be done in the standard way (see, e.g., Hull 2006) with daily rebalancing. Each basket option contract is written on 100,000 underlying shares and has a maturity of three months in trading days ($T = \frac{63}{252}$). The simulated asset price paths and option values are computed using the parameters given in Table 3.1 and the following correlation term structures for $t = \frac{1}{252}, \frac{2}{252}, \ldots, T$

Constant correlations:

$$\text{C1 to C9:}\quad \rho_t = \rho,$$

$$\text{where } \rho \in \{-0.9, -0.7, -0.5, -0.2, 0, 0.2, 0.5, 0.7, 0.9\}$$

Negative correlations jump upwards:

$$\begin{aligned}
\text{T1:}\quad \rho_t &= -0.9 + \left[0.9\left(t - \frac{1}{63}T\right)252\right]\mathbf{I}_{\left(\frac{1}{63}T<t\leq\frac{3}{63}T\right)} \\
&\quad + \left[1.8 - 0.9\left(t - \frac{3}{63}T\right)252\right]\mathbf{I}_{\left(\frac{3}{63}T<t\leq\frac{5}{63}T\right)} \\
\text{T2:}\quad \rho_t &= -0.9 + \left[0.9\left(t - \frac{59}{63}T\right)252\right]\mathbf{I}_{\left(\frac{59}{63}T<t\leq\frac{61}{63}T\right)} \\
&\quad + \left[1.8 - 0.9\left(t - \frac{61}{63}T\right)252\right]\mathbf{I}_{\left(\frac{61}{63}T<t\leq T\right)}
\end{aligned}$$

Positive correlations jump downwards:

$$\begin{aligned}
\text{T3:}\quad \rho_t &= 0.9 - 0.9\left[\left(t - \frac{1}{63}T\right)252\right]\mathbf{I}_{\left(\frac{1}{63}T<t\leq\frac{3}{63}T\right)} \\
&\quad - \left[1.8 - 0.9\left(t - \frac{3}{63}T\right)252\right]\mathbf{I}_{\left(\frac{3}{63}T<t\leq\frac{5}{63}T\right)} \\
\text{T4:}\quad \rho_t &= 0.9 - 0.9\left[\left(t - \frac{59}{63}T\right)252\right]\mathbf{I}_{\left(\frac{59}{63}T<t\leq\frac{61}{63}T\right)} \\
&\quad - \left[1.8 - 0.9\left(t - \frac{61}{63}T\right)252\right]\mathbf{I}_{\left(\frac{61}{63}T<t\leq T\right)}
\end{aligned}$$

Correlation regime shifts:

$$\text{T5:}\quad \rho_t = -0.9\mathbf{I}_{\left(t\leq\frac{30}{63}T\right)} + -0.5\mathbf{I}_{\left(\frac{30}{63}T<t\leq T\right)}$$

$$\text{T6:}\quad \rho_t = 0.9\mathbf{I}_{\left(t\leq\frac{30}{63}T\right)} + 0.5\mathbf{I}_{\left(\frac{30}{63}T<t\leq T\right)}$$

Affine correlation term structure:

$$\text{T7:}\quad \rho_t = -0.9 + \frac{1.8}{62}(252t - 1)$$

$$\text{T8:}\quad \rho_t = 0.9 - \frac{1.8}{62}(252t - 1)$$

Figure 2.2: True time-varying correlation term structures

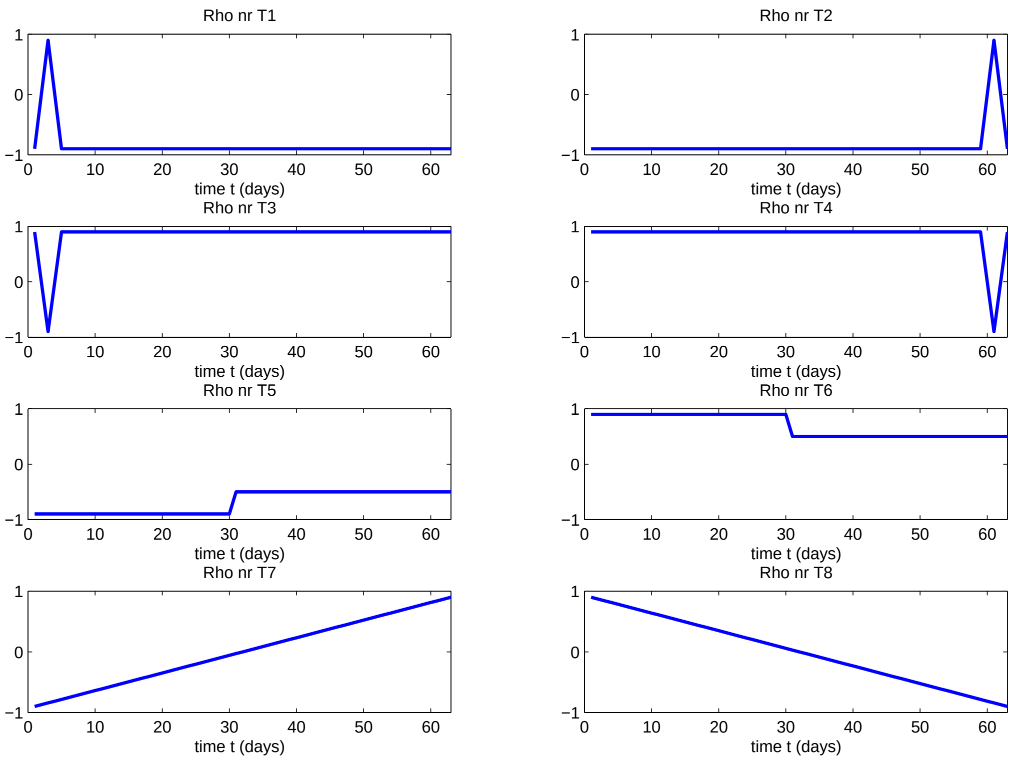

Figure 2.2 illustrates the time-varying correlation term structures, which can be interpreted as follows. The correlation term structures T1 and T2 correspond to a situation where the assets in the basket are usually well-diversified (i.e. negative correlations), but suddenly the correlations jump upwards at respectively at the start and the end of the option life. This could happen, for example, during a financial crisis, when all stocks plummet simultaneously and the correlations between these stocks suddenly jump upwards. Correlation specification T3 and T4 correspond to a situation where the correlation is highly positive, but suddenly jump downwards at respectively, the start and the end of the option life. T5 and T6 are regime-switching type of specification, where the correlations shift respectively upward or downward, but are locally constant before and after the shift. Finally, T7 corresponds to a gradually linear increase of the correlations from -0.9 to +0.9, often associated with a bear market in which stocks become more correlated over time. T8 represents a linear decrease in correlations, which

can be interpreted as a period in a bull market.

Eydeland and Wolyniec (2003) distinguish the instantaneous correlation ρ_t in (2.1) from the cumulative correlation defined by them as

$$\rho^*_{T,t} = \rho_{\ln S_{1,t} \ln S_{2,t}} = \frac{\mathbb{E}_t[\ln S_{1,T} \ln S_{2,T}] - \mathbb{E}_t[\ln S_{1,T}]\mathbb{E}_t[\ln S_{2,T}]}{\sqrt{\text{var}_t[\ln S_{1,T}]}\sqrt{\text{var}_t[\ln S_{2,T}]}} \tag{2.5}$$

$$= \frac{\int_t^T \sigma_{1,s}\sigma_{2,s}\rho_s ds}{\sqrt{\int_t^T \sigma^2_{1,s} ds}\sqrt{\int_t^T \sigma^2_{2,s} ds}}, \tag{2.6}$$

with the following properties

$$\lim_{t \to T} \rho^*_{T,t} = \rho_T \tag{2.7}$$

$$\text{if} \quad \sigma_{1,t} = \sigma_{2,t} = \sigma, \quad \forall t = 1, \cdots, T \quad \text{then} \quad \rho^*_{T,t} = \frac{\int_t^T \rho_s ds}{T - t} \tag{2.8}$$

According to Eydeland and Wolyniec (2003), the cumulative correlation $\rho^*_{T,t}$ is more important than the instantaneous correlation ρ_t for option pricing and hedging as well as for VaR computation. When the VaR horizon is very short, by (2.7) the cumulative correlation becomes close to the instantaneous correlation. By (2.8) the cumulative correlation boils down to the average correlation over time period $T - t$, if the volatilities of both assets are the same and constant over time. So, in this special case one needs only an estimate of the average value of the correlation term structure instead of information on the entire term structure to compute the VaR of the option. This is a great simplification and the results in the next section show to what extent this property holds for different time-varying correlation term structures.

Empirical correlations of asset returns are unobserved and have to be estimated using variance-covariance models, so it is very likely that the resulting correlations contain estimation errors. In this chapter the focus is on the effects of changes in the true correlation values instead of correlation estimation errors. For this purpose, the same parameter values will be used both to generate the asset price data and to value the basket options. Hence, the effects of changes in the true correlations can be isolated without bothering about estimation errors of correlations (for more details on correlation estimation errors, see Fengler and Schwendner (2004)).

Value-at-risk

In the current Basel II framework, banks should develop a more forward-looking approach with respect to risk management. The risks of the basket option position are assessed using the risk measure value-at-risk. Value-at-risk (VaR) is a widely used measure for quantifying potential losses of asset portfolios at a certain confidence level α (conventionally, 95 percent or 99 percent). Let X be the profit and loss realizations of the simulations, then VaR is defined by

$$\mathbb{P}(X \leq VaR^{\alpha}) = 1 - \alpha. \tag{2.9}$$

The main methods of VaR computation are the delta-normal method, historical simulation, and Monte Carlo simulation (see, e.g., Jorion 2001). To determine which method is suitable, it is important to look at the characteristics of the portfolio for which the VaR needs to be computed. The asset portfolio here consists of a basket option with (time-dependent) correlations between the underlying assets; therefore the risks involved in this asset position are highly nonlinear as well as time dependent. The Monte Carlo simulation method is used to compute the VaR in this study. The motivation for the choice of the Monte Carlo simulation method is, that it can take time dependency and nonlinear risk into account. Moreover, with respect to the requirements of Basel II, the computation of VaR can be easily adapted to reflect current market conditions by changing the underlying parameters such that these coincide with market parameters.

The alternative methods are more often used by market practitioners, since they are less computationally intensive than Monte Carlo simulation. However, the delta-normal method and the historical simulation method are less adequate for computing the VaR of the basket options for the following reasons. The delta-normal method is based on the assumption of a normally distributed portfolio, thus it is only suitable for portfolios involving linear risk in the underlyings. The historical simulation method is often used by market practitioners, since this method does not involve any distributional assumptions of the portfolio and is relatively fast to compute. Historical returns are used to represent potential future losses, and this implies that all information of the historical data is retained. However, the resulting VaR estimate can only reflect the type of losses which already have occurred in the historical data sample used for estimation. Hence, this approach is not in line with the forward-looking approach required in the Basel II framework. Moreover, due to the use of a relatively large number of historical observations to avoid small-sample bias, the most recent market movements cannot be easily captured in the VaR computation. For more details on the historical simulation method, see Pritsker (2006) and the references therein.

CVaR and the square-root-of-time rule

This study will assess the risks of the option portfolio using VaR, because this risk measure is widely used for regulatory purposes as specified by the Basel Committee on Banking Supervision. However, it is well known that VaR is only an indication of the loss corresponding to a confidence level α and over a certain time horizon (usually ten days); it does not give an indication of the size of that loss if the VaR is exceeded. Moreover, it is not a coherent risk measure (for more details, see Artzner et al 1999). Therefore, the conditional value-at-risk (CVaR) will also be computed, which satisfies the coherence properties. Let X be the profit and loss realizations of the simulations. The CVaR (also called expected shortfall) is defined for a given confidence level α as

$$CVaR^{\alpha} = \mathbb{E}[X|X \leq VaR^{\alpha}]. \tag{2.10}$$

The difference between the CVaR and the VaR can be expressed by the CVaR-to-VaR ratio. McNeil, Frey and Embrechts (2005) describe that for the CVaR-to-VaR ratio of the normal distribution it holds that $\lim_{\alpha \to 1} \frac{\text{CVaR}^{\alpha}}{\text{VaR}^{\alpha}} = 1$, whereas for the t-distribution this ratio will go to a value greater than 1 in the limit. Hence, for a heavy-tailed distribution the difference between the CVaR and VaR will be larger than for a normal distribution. The 1996 Amendment of the 1988 Basel Capital Accord requires a capital charge for market risk (see e.g. Jorion (2001) and Hull (2006) for more details). The market risk capital charge by the internal models approach can be computed as the maximum of the VaR of the previous day and the average VaR over the previous sixty days times a stress factor k. The stress factor k is determined by the local regulators and has at least an absolute value of 3. In practice, the square-root-of-time rule is often applied to compute the ten-day VaR and CVaR by multiplying the one-day VaR with $\sqrt{10}$, since the direct estimation of the ten-day VaR often requires too much historical data. This approach is valid for i.i.d. normally distributed observations; otherwise the square-root-of-time rule gives an approximation of the true VaR (see Jorion 2001).

2.4 Results

In this section, the results obtained for basket options with time-varying correlation term structures will be discussed using the benchmark of constant correlation term structures. To gain a better understanding of the results of time-varying correlations, the results obtained from constant correlations are first discussed.

2.4.1 Constant correlations

To illustrate how basket option prices of different moneyness levels change with respect to different constant correlation values between the underlyings, the option prices are plotted in Figure 2.3 and are calculated with the parameters in Table 3.1. Figure 2.3 shows the basket option prices that are computed for different moneyness levels using constant correlation values 0, ± 0.5 and ± 0.9 for the underlying assets. From this figure one can see that the differences between the option prices across correlation specifications are relatively pronounced for (near) ATM option prices. Moreover, the most striking observation from Figure 2.3 is that there is an asymmetric correlation effect on the ATM option prices: changes in negative correlations have a greater impact on the option prices than changes of the same magnitude in positive correlations (for example, the difference between the option prices for ρ equal to -0.9 and zero is greater than the difference between the prices of ρ equal to +0.9 and zero). This asymmetric correlation effect can also be seen for the OTM option prices, but to a lesser extent than for the ATM option prices. In Figure 2.3 the option prices converge to the same value beyond a certain moneyness level regardless of the correlation value for the underlying (see e.g. beyond S_0 of 100). However, the subsequent discussion of the simulation results shows that the risk of potential loss on the far ITM option over a certain time horizon does depend substantially on the correlation values of the underlying assets.

Figure 2.3: Basket option prices for different correlations for S0 from 90 to 110

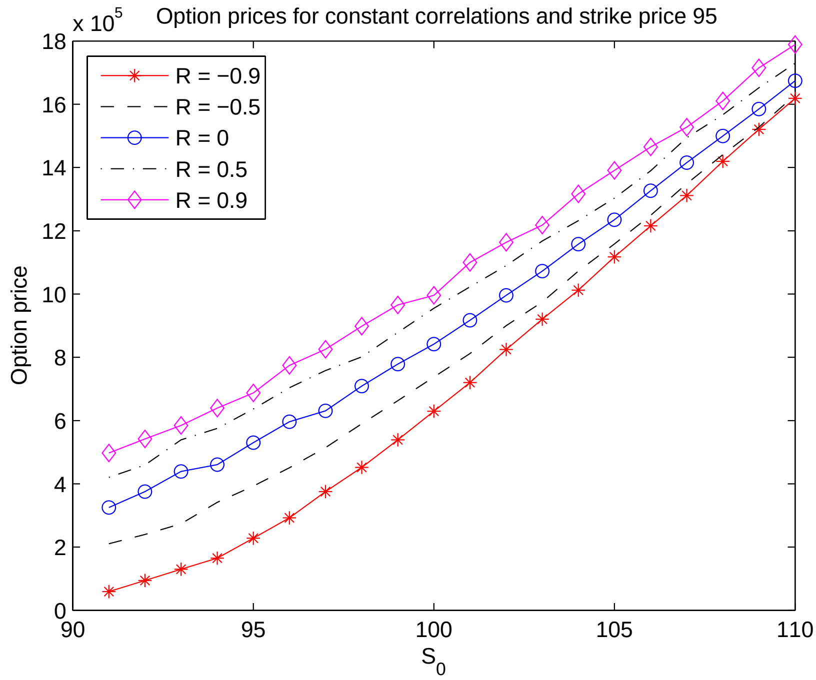

Asymmetric correlation effects

In this section the asymmetric effects of negative and positive correlation values on the option price as well as the risk measures VaR and CVaR of the basket option position will be discussed. The results of the simulation study for constant correlations are given in Table 2.2 and turn out to be highly dependent on the initial moneyness level of the basket option. To discuss these results in more detail, let $\bar{V}_0$ denote the average basket option price of the simulation sample at time 0 and define[3]

$$\begin{aligned} \delta \mathrm{VaR}^{\alpha}_{-} &= \mathrm{VaR}^{\alpha}(\rho = 0) - \mathrm{VaR}^{\alpha}(\rho = -0.9) \\ \delta \mathrm{VaR}^{\alpha}_{+} &= \mathrm{VaR}^{\alpha}(\rho = 0.9) - \mathrm{VaR}^{\alpha}(\rho = 0) \\ \delta \mathrm{VaR}^{\alpha}_{\mathrm{Total}} &= \delta \mathrm{VaR}^{\alpha}_{-} + \delta \mathrm{VaR}^{\alpha}_{+}. \end{aligned}$$

The results for the unhedged basket option are as follows. First, as the constant correlation ρ increases from -0.9 to 0.9 (respectively, C1 and C9), the average initial

[3]The VaR here is the ten-day VaR^{α} of the unhedged option position given in Table 2.2

option price $\bar{V}_0$ as well as the potential loss measured by VaR and CVaR increase for each moneyness level. For ITM options, this increase in the potential loss is unproportionally high relative to the increase in the average initial option price. For example, $\bar{V}_0^{ITM}$ of C1 and C9 are respectively $6.30 \cdot 10^5$ and $1.01 \cdot 10^6$, which is an increase with factor 1.61. The corresponding unhedged ten-day $\mathrm{VaR}^{0.95}$ estimates for C1 and C9 as given in Table 2.2 are $-2.68 \cdot 10^5$ to $-8.67 \cdot 10^5$ respectively, so the potential losses increase with a factor of 3.23 (this factor is even larger for $\mathrm{VaR}^{0.99}$ and CVaR). As the correlations change from -0.9 to +0.9, the $(1-\alpha)$-quantiles of the option position's profits and losses increase to more than three times larger than the $\bar{V}_0$ for ITM options. In contrast to this result, higher correlation values affect $\bar{V}_0^{ATM}$ option prices and the corresponding VaR and CVaR estimates more proportionally, where $\bar{V}_0^{ATM}$ increases with factor 3.08 and (C)VaR estimates increase with about factor 4 when going from correlation value -0.9 to +0.9.

Secondly, the largest part of the increase in potential losses for ITM and ATM options can be found by increasing the negative correlations. For example, from Table 2.2 the ITM $\delta\mathrm{VaR}_{-}^{0.95}$ has a value of $-3.92 \cdot 10^5$, which constitutes 65 percent of the total change $\delta\mathrm{VaR}_{\mathrm{Total}}^{0.95}$ of $-5.98 \cdot 10^5$. For the ATM and OTM option, this percentage can be obtained similarly and is 63 percent of the $\delta\mathrm{VaR}_{\mathrm{Total}}^{0.95}$. The relatively high sensitivity of basket options on well-diversified baskets (negative correlations) with respect to correlation changes seems counterintuitive. It is well known that well-diversified portfolios are less risky than portfolios with highly correlated assets; therefore investors might expect that options on a well-diversified portfolio are also less risky than options on highly correlated underlying assets. Although the results in Table 2.2 show that the absolute *values* of the (C)VaR of the basket option increase with higher correlations, the *changes* in these (C)VaR values are relatively more subject to correlation risk for a negatively correlated portfolio. This observation is important for investors, because they may not be aware of this increased correlation risk, when buying basket options on well-diversified ('safe') basket of assets. Another implication of this result is that VaR estimation of basket options on well-diversified baskets is relatively more prone to model risk, since in practice correlations have to be estimated under the assumption of a correlation model (e.g. RiskMetrics^{TM}), which is in line with the results of Skintzi, Skiadopoulos and Refenes (2005). They have found that the VaR measure becomes relatively more sensitive to correlation estimation errors with decreasing true correlations for linear portfolios as well as option portfolios containing plain-vanilla European options written on correlated underlying assets.

Finally, the option prices and risk measures of OTM options vary even more with the correlation of the underlying assets than those of ITM and ATM options. The $\bar{V}_0$ prices for ρ of -0.9 and +0.9 are respectively $5.45 \cdot 10^4$ and $5.23 \cdot 10^5$ and the corresponding $\text{VaR}^{0.95}$ estimates are $-7.47 \cdot 10^4$ and $-6.11 \cdot 10^5$. So, the difference in option price and risk is extremely large across different correlation values for OTM basket option position.

Delta-hedged portfolios

In practice, option positions of financial institutions are usually hedged to decrease the exposure to the full risks of the option position. Therefore, it is perhaps even more important to consider the effects of correlations on daily delta-hedged option positions. A delta-hedged option position accounts for the risk of changes in the price of the underlying assets. The estimates of the risk measures are given in Table 2.2. The effectiveness of the delta hedge in reducing the risk of the option position will be measured by the no-hedge-to-hedge ratio of the risk measures given in Table 2.3. If this ratio is very high, the risk of potential loss is reduced much better by hedging than in case of a low ratio. The results are very different across moneyness levels.

For ITM options, daily delta hedging reduces the risk by a factor of 23 or more for the highly negative correlation specification C1 for all risk measures. The hedge effectiveness reduces for higher correlations, but it is still considerable for positive correlations, as the no-hedge-to-hedge ratio of the risk measures has a value of more than 16 for different confidence levels.

The delta hedge effectiveness deteriorates for ATM and OTM options for all values of correlations. Still, the risk reduction is substantial compared to the unhedged ATM and OTM option position in most cases, and the no-hedge-to-hedge ratios vary from around 7 to 15. In contrast to the result for ITM options, the reduction in risk achieved by delta hedging is the largest for highly positive correlations (especially for OTM options). Overall, the delta hedge is most effective for ITM options but still substantial for ATM and OTM options.

Table 2.2: VaR and CVaR results for ITM options with constant correlations. The abbreviations in the table denote the following: NH = no hedging, HE = daily delta hedging, $\sqrt{10}$NH1 = $\sqrt{10}\text{VaR}_{1-day}$ (no hedging), $\sqrt{10}$HE1 = $\sqrt{10}\text{VaR}_{1-day}$(daily delta hedging)

	$\text{VaR}^{0.99}$	$\text{VaR}^{0.95}$	$\text{CVaR}^{0.99}$	$\text{CVaR}^{0.95}$	$\frac{\text{CVaR}^{0.99}}{\text{VaR}^{0.99}}$
ITM C1 NH	$-3.66 \cdot 10^5$	$-2.68 \cdot 10^5$	$-4.24 \cdot 10^5$	$-3.29 \cdot 10^5$	1.16
ITM C1 HE	$-1.54 \cdot 10^4$	$-1.08 \cdot 10^4$	$-1.67 \cdot 10^4$	$-1.35 \cdot 10^4$	1.09
ITM C1 $\sqrt{10}$NH1	$-3.51 \cdot 10^5$	$-2.46 \cdot 10^5$	$-3.93 \cdot 10^5$	$-3.11 \cdot 10^5$	1.12
ITM C1 $\sqrt{10}$HE1	$-4.35 \cdot 10^4$	$-3.08 \cdot 10^4$	$-4.95 \cdot 10^4$	$-3.90 \cdot 10^4$	1.14
ITM C2 NH	$-5.60 \cdot 10^5$	$-4.14 \cdot 10^5$	$-6.63 \cdot 10^5$	$-5.10 \cdot 10^5$	1.19
ITM C2 HE	$-2.52 \cdot 10^4$	$-1.69 \cdot 10^4$	$-3.10 \cdot 10^4$	$-2.21 \cdot 10^4$	1.23
ITM C2 $\sqrt{10}$NH1	$-5.40 \cdot 10^5$	$-3.68 \cdot 10^5$	$-6.23 \cdot 10^5$	$-4.75 \cdot 10^5$	1.15
ITM C2 $\sqrt{10}$HE1	$-5.10 \cdot 10^4$	$-3.53 \cdot 10^4$	$-5.71 \cdot 10^4$	$-4.45 \cdot 10^4$	1.12
ITM C3 NH	$-6.98 \cdot 10^5$	$-5.06 \cdot 10^5$	$-8.11 \cdot 10^5$	$-6.28 \cdot 10^5$	1.16
ITM C3 HE	$-3.50 \cdot 10^4$	$-2.26 \cdot 10^4$	$-4.38 \cdot 10^4$	$-3.04 \cdot 10^4$	1.25
ITM C3 $\sqrt{10}$NH1	$-6.77 \cdot 10^5$	$-4.44 \cdot 10^5$	$-7.74 \cdot 10^5$	$-5.76 \cdot 10^5$	1.14
ITM C3 $\sqrt{10}$HE1	$-6.30 \cdot 10^4$	$-4.37 \cdot 10^4$	$-7.22 \cdot 10^4$	$-5.54 \cdot 10^4$	1.15
ITM C4 NH	$-8.48 \cdot 10^5$	$-6.07 \cdot 10^5$	$-9.86 \cdot 10^5$	$-7.67 \cdot 10^5$	1.16
ITM C4 HE	$-4.84 \cdot 10^4$	$-2.87 \cdot 10^4$	$-5.85 \cdot 10^4$	$-4.04 \cdot 10^4$	1.21
ITM C4 $\sqrt{10}$NH1	$-7.87 \cdot 10^5$	$-5.34 \cdot 10^5$	$-9.40 \cdot 10^5$	$-6.95 \cdot 10^5$	1.20
ITM C4 $\sqrt{10}$HE1	$-8.46 \cdot 10^4$	$-5.45 \cdot 10^4$	$-9.49 \cdot 10^4$	$-7.04 \cdot 10^4$	1.12
ITM C5 NH	$-9.42 \cdot 10^5$	$-6.60 \cdot 10^5$	$-1.10 \cdot 10^6$	$-8.43 \cdot 10^5$	1.17
ITM C5 HE	$-5.45 \cdot 10^4$	$-3.37 \cdot 10^4$	$-6.70 \cdot 10^4$	$-4.60 \cdot 10^4$	1.23
ITM C5 $\sqrt{10}$NH1	$-8.45 \cdot 10^5$	$-5.90 \cdot 10^5$	$-1.03 \cdot 10^6$	$-7.61 \cdot 10^5$	1.22
ITM C5 $\sqrt{10}$HE1	$-9.50 \cdot 10^4$	$-6.15 \cdot 10^4$	$-1.09 \cdot 10^5$	$-8.04 \cdot 10^4$	1.15
ITM C6 NH	$-1.05 \cdot 10^6$	$-7.05 \cdot 10^5$	$-1.20 \cdot 10^6$	$-9.12 \cdot 10^5$	1.15
ITM C6 HE	$-5.96 \cdot 10^4$	$-3.72 \cdot 10^4$	$-7.25 \cdot 10^4$	$-5.03 \cdot 10^4$	1.22
ITM C6 $\sqrt{10}$NH1	$-9.04 \cdot 10^5$	$-6.53 \cdot 10^5$	$-1.11 \cdot 10^6$	$-8.21 \cdot 10^5$	1.23
ITM C6 $\sqrt{10}$HE1	$-1.04 \cdot 10^5$	$-6.87 \cdot 10^4$	$-1.21 \cdot 10^5$	$-8.98 \cdot 10^4$	1.17
ITM C7 NH	$-1.19 \cdot 10^6$	$-7.60 \cdot 10^5$	$-1.35 \cdot 10^6$	$-1.01 \cdot 10^6$	1.13
ITM C7 HE	$-6.36 \cdot 10^4$	$-4.21 \cdot 10^4$	$-7.83 \cdot 10^4$	$-5.57 \cdot 10^4$	1.23
ITM C7 $\sqrt{10}$NH1	$-1.04 \cdot 10^6$	$-7.33 \cdot 10^5$	$-1.22 \cdot 10^6$	$-9.08 \cdot 10^5$	1.17
ITM C7 $\sqrt{10}$HE1	$-1.19 \cdot 10^5$	$-8.09 \cdot 10^4$	$-1.37 \cdot 10^5$	$-1.03 \cdot 10^5$	1.15
ITM C8 NH	$-1.21 \cdot 10^6$	$-8.12 \cdot 10^5$	$-1.43 \cdot 10^6$	$-1.07 \cdot 10^6$	1.19
ITM C8 HE	$-6.80 \cdot 10^4$	$-4.44 \cdot 10^4$	$-8.16 \cdot 10^4$	$-5.92 \cdot 10^4$	1.20
ITM C8 $\sqrt{10}$NH1	$-1.09 \cdot 10^6$	$-7.67 \cdot 10^5$	$-1.28 \cdot 10^6$	$-9.71 \cdot 10^5$	1.17
ITM C8 $\sqrt{10}$HE1	$-1.25 \cdot 10^5$	$-8.57 \cdot 10^4$	$-1.49 \cdot 10^5$	$-1.11 \cdot 10^5$	1.19
ITM C9 NH	$-1.30 \cdot 10^6$	$-8.67 \cdot 10^5$	$-1.51 \cdot 10^6$	$-1.14 \cdot 10^6$	1.16
ITM C9 HE	$-7.16 \cdot 10^4$	$-4.68 \cdot 10^4$	$-8.43 \cdot 10^4$	$-6.23 \cdot 10^4$	1.18
ITM C9 $\sqrt{10}$NH1	$-1.16 \cdot 10^6$	$-8.10 \cdot 10^5$	$-1.32 \cdot 10^6$	$-1.04 \cdot 10^6$	1.14
ITM C9 $\sqrt{10}$HE1	$-1.40 \cdot 10^5$	$-9.25 \cdot 10^4$	$-1.61 \cdot 10^5$	$-1.20 \cdot 10^5$	1.15

Risk measures and square-root-of-time rule

This section discusses the difference between the VaR and CVaR risk measures in this simulation experiment as well as the performance of the widely used square-root-of-time rule. Overall, the risk of potential loss increases for a higher value of the underlying

Table 2.3: Risk measure ratios for ITM options with constant correlations. The abbreviations in the table denote the following: NH = no hedging, HE = daily delta hedging, $\sqrt{10}$NH1 = $\sqrt{10}\text{VaR}_{1-day}$ (no hedging), $\sqrt{10}$HE1 = $\sqrt{10}\text{VaR}_{1-day}$(daily delta hedging)

	$\text{VaR}^{0.99}$	$\text{VaR}^{0.95}$	$\text{CVaR}^{0.99}$	$\text{CVaR}^{0.95}$
ITM C1 NH/HE	23.86	24.89	25.33	24.48
ITM C1 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.08	8.00	7.93	7.96
ITM C1 NH/$\sqrt{10}$NH1	1.04	1.09	1.08	1.06
ITM C1 HE/$\sqrt{10}$HE1	0.35	0.35	0.34	0.35
ITM C2 NH/HE	22.18	24.50	21.38	23.10
ITM C2 $\sqrt{10}$NH1/$\sqrt{10}$HE1	10.59	10.44	10.90	10.68
ITM C2 NH/$\sqrt{10}$NH1	1.04	1.13	1.06	1.07
ITM C2 HE/$\sqrt{10}$HE1	0.50	0.48	0.54	0.50
ITM C3 NH/HE	19.96	22.41	18.53	20.69
ITM C3 $\sqrt{10}$NH1/$\sqrt{10}$HE1	10.75	10.14	10.72	10.40
ITM C3 NH/$\sqrt{10}$NH1	1.03	1.14	1.05	1.09
ITM C3 HE/$\sqrt{10}$HE1	0.56	0.52	0.61	0.55
ITM C4 NH/HE	17.51	21.12	16.86	18.96
ITM C4 $\sqrt{10}$NH1/$\sqrt{10}$HE1	9.30	9.80	9.91	9.88
ITM C4 NH/$\sqrt{10}$NH1	1.08	1.14	1.05	1.10
ITM C4 HE/$\sqrt{10}$HE1	0.57	0.53	0.62	0.58
ITM C5 NH/HE	17.28	19.60	16.41	18.35
ITM C5 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.89	9.59	9.41	9.46
ITM C5 NH/$\sqrt{10}$NH1	1.12	1.12	1.07	1.11
ITM C5 HE/$\sqrt{10}$HE1	0.57	0.55	0.61	0.57
ITM C6 NH/HE	17.57	18.93	16.60	18.14
ITM C6 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.73	9.51	9.20	9.15
ITM C6 NH/$\sqrt{10}$NH1	1.16	1.08	1.08	1.11
ITM C6 HE/$\sqrt{10}$HE1	0.58	0.54	0.60	0.56
ITM C7 NH/HE	18.74	18.07	17.24	18.14
ITM C7 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.72	9.06	8.90	8.83
ITM C7 NH/$\sqrt{10}$NH1	1.15	1.04	1.11	1.11
ITM C7 HE/$\sqrt{10}$HE1	0.53	0.52	0.57	0.54
ITM C8 NH/HE	17.72	18.28	17.59	18.15
ITM C8 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.73	8.96	8.59	8.75
ITM C8 NH/$\sqrt{10}$NH1	1.10	1.06	1.12	1.11
ITM C8 HE/$\sqrt{10}$HE1	0.54	0.52	0.55	0.53
ITM C9 NH/HE	18.13	18.52	17.87	18.25
ITM C9 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.31	8.75	8.23	8.67
ITM C9 NH/$\sqrt{10}$NH1	1.12	1.07	1.14	1.10
ITM C9 HE/$\sqrt{10}$HE1	0.51	0.51	0.53	0.52

constant correlation (see the VaR and CVaR estimates in Table 2.2). From Table 2.3 one can see that the $\text{CVaR}^{0.99}$-to-$\text{VaR}^{0.99}$ ratios of the ITM option are larger for the hedged option position than for the unhedged position in most cases, thus implying that the hedged profit and loss distribution has relatively heavier tails. However, for the ATM and OTM option the profit and loss distribution of the unhedged option position exhibits heavier tails compared with the hedged option position. Define γ as

$$\gamma = \frac{\text{ten-day VaR}^{0.99}}{\sqrt{10} \cdot \text{one-day VaR}^{0.99}}, \tag{2.11}$$

where γ is used to measure the performance of the square-root-of-time rule. If γ is larger than one, the unhedged ten-day VaR is underestimated when applying the square-root-of-time rule. One would expect that the square-root-of-time rule performs worse for cases with high CVaR$^{0.99}$-to-VaR$^{0.99}$ ratios, since the validity of this rule critically depends on the assumption of i.i.d.-normal observations. This expectation is indeed confirmed by the results, since the CVaR$^{0.99}$-to-VaR$^{0.99}$ ratios for the unhedged OTM option are relatively high compared to the corresponding values for the unhedged ATM option, and this also holds for the underestimation of the ten-day VaR$^{0.99}$. A notable exception is the unhedged ITM option position, where the square-root-of-time rule performs quite well. For the hedged option position, there is overestimation of the ten-day VaR$^{0.99}$ by the square-root-of-time rule, since daily delta hedging has an offsetting effect on the profit and loss changes of the option position. The no-hedge-to-hedge ratios are quite high, with a minimum value of about 7. This overestimation declines as the difference of the hedge effectiveness between the one-day and ten-day horizon diminishes, i.e., as the difference in the no-hedge-to-hedge ratios decreases.

So, in most cases, the square-root-of-time rule leads to either underestimation or overestimation of the ten-day VaR$^{0.99}$. However, none of the γ ratios exceed the minimum absolute stress factor value of 3. Hence, for the computation of market risk capital charge, the square-root-of-time rule used in combination with the stress factor k of at least 3 is reasonable here for basket options in case of constant correlations.

2.4.2 Time-varying correlations

In this section the results of time-varying correlation term structures containing jumps, regime switches, and affine term structures will be discussed for the basket option. By (2.8) the computation of the VaR for basket options with time-varying correlation term structures can be greatly simplified: one only needs to know an estimate of the average correlation over the option life instead of information on the entire future correlation term structure. The results in Table 2.4 show that property (2.8) does not hold for the basket option with different correlation term structures. If the correlations are relatively high during the ten-day horizon over which the VaR is computed, then the potential loss will be higher than indicated by the VaR estimated using the average correlation over the life of the option (and vice versa). This result is in line with Eydeland and Wolyniec (2003), where they explain that for a short VaR horizon of ten days, the instantaneous correlation becomes more important. Hence, there are differences in the results depending on the specific type of time variation, even though the average correlation over the option life is the same, and this will be discussed in more detail later. This result implies that, in practice, the cumulative correlation might not be as relevant for VaR computations as expected. The reduction in risks achieved by daily delta hedging is often less for time-varying correlation term structures than for the constant (cumulative) correlations, but the risk reduction is still substantial for ITM or ATM options and, to a lesser, extent for OTM options. The results for the hedged option are highly dependent on the hedge effectiveness. The following discussion deals with the differences in the specific time variation of the correlation term structures T1 to T8.

Correlation jumps

The correlation term structures T1 and T2 are initially highly negative and jump upward, respectively, at the start and the end of the option life. First, the time at which the correlation jump occurs is important for the VaR and CVaR estimates of the unhedged option, but has an even larger impact on the effectiveness of the delta hedge. When no hedging is applied, a jump at the start of the option life is more than twice as risky as a jump near expiration for negative correlations, see Table 2.4. For the delta-hedged option position this result is even stronger, since the risk measures indicate that the potential loss for a jump at the start is more than four times larger than for a jump near expiration. These results are robust for different moneyness levels. The average correlation for T1 and T2 is -0.84. The risk measures of only T2 give results that are comparable to those of C1 (constant correlation value of -0.9), but the risks for T1 are much higher than for C1. The no-hedge-to-hedge ratios of the risk measures of T2 in Table 2.5 are twice as large as for T1. So, the time of the temporary jump upward in the correlation has a large impact on the effectiveness of the delta hedge across all moneyness levels.

When the correlations are initially positive, a jump downward at the start (T3) is less risky than a jump near expiration (T4) for the unhedged basket option, regardless of the moneyness level. The hedge effectiveness of a jump near expiration is greater than a jump at the start for positive correlations. So, T2 and T4 have a relatively better hedging performance compared with, respectively, T1 and T3, as shown by the no-hedge-to-hedge ratios in Table 2.5. Hence, delta hedging of the basket option is more effective in reducing the risk of potential loss when the correlation jump does not occur within the VaR horizon of ten days. The mean and standard deviation of the initial option prices are nearly the same for T1 and T2, as well as for T3 and T4. Therefore, stress testing is quite important, since the valuation of the option does not show any of these differences in the potential loss of the position.

Correlation regime switches

The correlation term structure T5 has initially highly negative (constant) correlations but shifts upwards in the second half of the life of the option. For the unhedged ITM and ATM option, a jump near expiration (T2) is riskier than a correlation shift (T5), although the jump only lasts for five days (see the risk measures in Table 2.4).

Due to the relatively high hedge effectiveness of T2 for the ITM option compared with T5 (i.e., high no-hedge-to-hedge ratios), the (C)VaR estimates show that the potential loss of a jump near expiration is lower than a correlation shift for the hedged ITM option. For the hedged ATM option, the no-hedge-to-hedge ratios of T2 and T5 are comparable, and a jump near expiration still has a higher potential loss than a correlation shift for negative correlations. Regarding the OTM option, a correlation shift is riskier than a jump near expiration when no hedging is applied, and this also holds for the hedged option according to the (C)VaR estimates at the 5 percent significance level.

The risk measures do not differ largely between T4 and T6, which are correlation term structures with initially positive correlations and, respectively, a correlation jump and shift downward. The hedge effectiveness is much larger for T6 than for T5 in case of ATM and OTM options in terms of higher no-hedge-to-hedge ratios of risk measures.

Affine correlation term structures

The linear increase and decrease of, respectively, correlations T7 and T8 have cumulative correlation of zero, which corresponds to C5. For C5, T7, and T8, the initial option prices have very similar mean and standard deviations. The unhedged (C)VaR results of T7 are much lower than those of C5, but T8 shows relatively higher risk of potential loss. However, the hedge effectiveness of T8 is relatively higher than for T7, as shown by the higher no-hedge-to-hedge ratios for T8. During bull markets, the correlations could decline; during bear markets or crisis, correlations often increase. Although the unhedged basket option is riskier for declining correlations (e.g., bull markets) in terms of higher potential loss estimates, the daily delta hedge is much more effective in reducing risk than for increasing correlations (e.g., bear markets).

Risk measures and square-root-of-time rule

When comparing the ten-day VaR to the ten-day CVaR, the results in Table 2.4 show that the difference in the risk measures can be quite large. For example, the potential loss indicated by the $\text{CVaR}^{0.99}$ is more than 40 percent larger than by the $\text{VaR}^{0.99}$ for the hedged ITM option with correlation term structure T1. For the other term structures the difference is less pronounced but can still be substantial, especially at the 5 percent significance level. Therefore, it is advisable to look at both the VaR and CVaR estimates, since the CVaR could provide additional information on the risk exposure.

Many market practitioners use the square-root-of-time rule. Previously, the VaR and CVaR results for constant correlations have shown that the square-root-of-time rule provides reasonable estimates for the unhedged basket option. To my knowledge, there is not much evidence in the literature on how the square-root-of-time rule for VaR and CVaR performs for basket options if the correlation term structure is time varying. In this chapter, simulation results show that the square-root-of-time rule (C)VaR estimates substantially underestimate the ten-day (C)VaR for the unhedged option with time-varying correlation term structures as shown in Table 2.4. This underestimation is relatively severe for the cases T1 and T7, where the negative correlation between the underlying assets suddenly increases (as is often observed in financial crises). If the correlations at the first day of the VaR horizon are highly negative, then the VaR given by the square-root-of-time rule is computed based on this negative value and

Table 2.4: VaR and CVaR results for ITM options with time-varying correlations. The abbreviations in the table denote the following: NH = no hedging, HE = daily delta hedging, $\sqrt{10}$NH1 = $\sqrt{10}\text{VaR}_{1-day}$ (no hedging), $\sqrt{10}$HE1 = $\sqrt{10}\text{VaR}_{1-day}$(daily delta hedging)

	$\text{VaR}^{0.99}$	$\text{VaR}^{0.95}$	$\text{CVaR}^{0.99}$	$\text{CVaR}^{0.95}$	$\frac{\text{CVaR}^{0.99}}{\text{VaR}^{0.99}}$
ITM T1 NH	$-8.39 \cdot 10^5$	$-5.55 \cdot 10^5$	$-9.27 \cdot 10^5$	$-7.17 \cdot 10^5$	1.11
ITM T1 HE	$-5.96 \cdot 10^4$	$-3.19 \cdot 10^4$	$-8.58 \cdot 10^4$	$-5.12 \cdot 10^4$	1.44
ITM T1 $\sqrt{10}$NH1	$-3.40 \cdot 10^5$	$-2.35 \cdot 10^5$	$-3.78 \cdot 10^5$	$-2.97 \cdot 10^5$	1.11
ITM T1 $\sqrt{10}$HE1	$-4.29 \cdot 10^4$	$-3.24 \cdot 10^4$	$-4.84 \cdot 10^4$	$-3.89 \cdot 10^4$	1.13
ITM T2 NH	$-3.55 \cdot 10^5$	$-2.59 \cdot 10^5$	$-4.12 \cdot 10^5$	$-3.18 \cdot 10^5$	1.16
ITM T2 HE	$-1.49 \cdot 10^4$	$-1.10 \cdot 10^4$	$-1.68 \cdot 10^4$	$-1.35 \cdot 10^4$	1.12
ITM T2 $\sqrt{10}$NH1	$-3.37 \cdot 10^5$	$-2.36 \cdot 10^5$	$-3.78 \cdot 10^5$	$-2.98 \cdot 10^5$	1.12
ITM T2 $\sqrt{10}$HE1	$-4.52 \cdot 10^4$	$-3.08 \cdot 10^4$	$-5.02 \cdot 10^4$	$-3.96 \cdot 10^4$	1.11
ITM T3 NH	$-1.15 \cdot 10^6$	$-7.36 \cdot 10^5$	$-1.47 \cdot 10^6$	$-9.98 \cdot 10^5$	1.28
ITM T3 HE	$-6.82 \cdot 10^4$	$-4.24 \cdot 10^4$	$-8.08 \cdot 10^4$	$-5.80 \cdot 10^4$	1.19
ITM T3 $\sqrt{10}$NH1	$-1.16 \cdot 10^6$	$-8.11 \cdot 10^5$	$-1.33 \cdot 10^6$	$-1.04 \cdot 10^6$	1.14
ITM T3 $\sqrt{10}$HE1	$-1.38 \cdot 10^5$	$-9.13 \cdot 10^4$	$-1.60 \cdot 10^5$	$-1.19 \cdot 10^5$	1.16
ITM T4 NH	$-1.30 \cdot 10^6$	$-8.72 \cdot 10^5$	$-1.51 \cdot 10^6$	$-1.14 \cdot 10^6$	1.16
ITM T4 HE	$-7.08 \cdot 10^4$	$-4.72 \cdot 10^4$	$-8.50 \cdot 10^4$	$-6.19 \cdot 10^4$	1.20
ITM T4 $\sqrt{10}$NH1	$-1.16 \cdot 10^6$	$-8.07 \cdot 10^5$	$-1.32 \cdot 10^6$	$-1.04 \cdot 10^6$	1.14
ITM T4 $\sqrt{10}$HE1	$-1.32 \cdot 10^5$	$-8.98 \cdot 10^4$	$-1.55 \cdot 10^5$	$-1.18 \cdot 10^5$	1.18
ITM T5 NH	$-3.30 \cdot 10^5$	$-2.40 \cdot 10^5$	$-3.86 \cdot 10^5$	$-2.97 \cdot 10^5$	1.17
ITM T5 HE	$-1.66 \cdot 10^4$	$-1.22 \cdot 10^4$	$-1.83 \cdot 10^4$	$-1.49 \cdot 10^4$	1.10
ITM T5 $\sqrt{10}$NH1	$-3.15 \cdot 10^5$	$-2.15 \cdot 10^5$	$-3.52 \cdot 10^5$	$-2.73 \cdot 10^5$	1.12
ITM T5 $\sqrt{10}$HE1	$-4.82 \cdot 10^4$	$-3.24 \cdot 10^4$	$-5.42 \cdot 10^4$	$-4.16 \cdot 10^4$	1.12
ITM T6 NH	$-1.32 \cdot 10^6$	$-8.79 \cdot 10^5$	$-1.53 \cdot 10^6$	$-1.15 \cdot 10^6$	1.16
ITM T6 HE	$-7.36 \cdot 10^4$	$-4.78 \cdot 10^4$	$-8.60 \cdot 10^4$	$-6.32 \cdot 10^4$	1.17
ITM T6 $\sqrt{10}$NH1	$-1.18 \cdot 10^6$	$-8.17 \cdot 10^5$	$-1.34 \cdot 10^6$	$-1.05 \cdot 10^6$	1.14
ITM T6 $\sqrt{10}$HE1	$-1.29 \cdot 10^5$	$-8.60 \cdot 10^4$	$-1.50 \cdot 10^5$	$-1.13 \cdot 10^5$	1.17
ITM T7 NH	$-4.25 \cdot 10^5$	$-3.15 \cdot 10^5$	$-4.97 \cdot 10^5$	$-3.79 \cdot 10^5$	1.17
ITM T7 HE	$-2.89 \cdot 10^4$	$-2.06 \cdot 10^4$	$-3.31 \cdot 10^4$	$-2.55 \cdot 10^4$	1.15
ITM T7 $\sqrt{10}$NH1	$-2.86 \cdot 10^5$	$-1.97 \cdot 10^5$	$-3.24 \cdot 10^5$	$-2.47 \cdot 10^5$	1.13
ITM T7 $\sqrt{10}$HE1	$-8.14 \cdot 10^4$	$-5.63 \cdot 10^4$	$-9.15 \cdot 10^4$	$-7.12 \cdot 10^4$	1.12
ITM T8 NH	$-1.32 \cdot 10^6$	$-9.24 \cdot 10^5$	$-1.57 \cdot 10^6$	$-1.18 \cdot 10^6$	1.19
ITM T8 HE	$-7.90 \cdot 10^4$	$-5.09 \cdot 10^4$	$-1.02 \cdot 10^5$	$-7.00 \cdot 10^4$	1.29
ITM T8 $\sqrt{10}$NH1	$-1.25 \cdot 10^6$	$-8.54 \cdot 10^5$	$-1.40 \cdot 10^6$	$-1.10 \cdot 10^6$	1.12
ITM T8 $\sqrt{10}$HE1	$-1.24 \cdot 10^5$	$-7.45 \cdot 10^4$	$-1.55 \cdot 10^5$	$-1.04 \cdot 10^5$	1.25

thus neglects the fact that the correlations increase over time.

The ratio γ of the ten-day $\text{VaR}^{0.99}$ to the square-root-of-time rule estimate is given in Table 2.3 for the constant correlations and in Table 2.5 for time-varying correlation term structures. The minimum of the stress factor k is 3 and in Table 2.3 it is shown that the γ from our VaR results for the constant correlation specification do not exceed 3 for all constant correlation values and moneyness levels. The underestimation of the unhedged ten-day VaR by the square-root-of-time rule stays well below a factor of 1.5

for constant correlations. The results in Tables 2.5 show that the underestimation is much more severe for time-varying correlations, where the unhedged ten-day $\mathrm{VaR}^{0.99}$ for T1 is underestimated with a factor of 2.47 for the unhedged ITM option and increases to 4.08 for the OTM option, thus even violating the minimum stress factor k of 3.

The performance of the square-root-of-time rule is worse for the hedged option as can be seen in Table 2.5. Depending on the difference between the hedge effectiveness for the one-day and ten-day horizon, the square-root-of-time rule could lead to either severe underestimation or overestimation of the ten-day (C)VaR. For example, the no-hedge-to-hedge ratios for the ten-day $\mathrm{VaR}^{0.99}$ and the scaled one-day $\mathrm{VaR}^{0.99}$ are, respectively, 14.08 and 7.93 for T1 in Table 2.5. The T1 results for the unhedged ITM option show an underestimation of the ten-day $\mathrm{VaR}^{0.99}$ by the square-root-of-time rule, and this also holds for the ten-day $\mathrm{VaR}^{0.99}$ of the hedged option. For T7, the no-hedge-to-hedge ratios of the ten-day $\mathrm{VaR}^{0.99}$ and the scaled one-day $\mathrm{VaR}^{0.99}$ are, respectively, 14.71 and 3.51. Hence, the hedge effectiveness for the T7 scaled one-day $\mathrm{VaR}^{0.99}$ is relatively much lower than for the ten-day $\mathrm{VaR}^{0.99}$ in the case of T1. The hedged ten-day $\mathrm{VaR}^{0.99}$ shows that the risk of potential loss for T1 is much larger than for T7, but the relatively low hedge effectiveness for the one-day horizon of T7 leads to the result that the scaled one-day $\mathrm{VaR}^{0.99}$ of T7 is twice as large as the same risk measure for T1. Thus, the square-root-of-time rule leads to an overestimation of the T7 ten-day (C)VaR for the hedged option due to this large difference in hedge effectiveness between the one-day and ten-day horizon. The ten-day $\mathrm{VaR}^{0.99}$ is also overestimated by the square-root-of-time rule for many other time-varying correlation term structures for the hedged option.

Since the performance of the square-root-of-time rule for hedged options is very sensitive to the hedge effectiveness, the minimum stress factor is likely to be violated in situations where hedging becomes relatively difficult — for example, for options that are out-of-the money or near expiration. Hence, the square-root-of-time rule applied to a one-day (C)VaR to obtain a ten-day (C)VaR appears to be inadequate for time-varying correlation term structures, even in this simplified and ideal simulation environment.

Table 2.5: Risk measure ratios for ITM options with time-varying correlations The abbreviations in the table denote the following: NH = no hedging, HE = daily delta hedging, $\sqrt{10}$NH1 = $\sqrt{10}\text{VaR}_{1-day}$ (no hedging), $\sqrt{10}$HE1 = $\sqrt{10}\text{VaR}_{1-day}$(daily delta hedging)

	$\text{VaR}^{0.99}$	$\text{VaR}^{0.95}$	$\text{CVaR}^{0.99}$	$\text{CVaR}^{0.95}$
ITM T1 NH/HE	14.08	17.43	10.81	13.99
ITM T1 $\sqrt{10}$NH1/$\sqrt{10}$HE1	7.93	7.27	7.80	7.62
ITM T1 NH/$\sqrt{10}$NH1	2.47	2.36	2.46	2.42
ITM T1 HE/$\sqrt{10}$HE1	1.39	0.99	1.77	1.32
ITM T2 NH/HE	23.75	23.44	24.53	23.57
ITM T2 $\sqrt{10}$NH1/$\sqrt{10}$HE1	7.47	7.64	7.52	7.52
ITM T2 NH/$\sqrt{10}$NH1	1.05	1.10	1.09	1.07
ITM T2 HE/$\sqrt{10}$HE1	0.33	0.36	0.34	0.34
ITM T3 NH/HE	16.85	17.36	18.16	17.23
ITM T3 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.46	8.88	8.30	8.78
ITM T3 NH/$\sqrt{10}$NH1	0.99	0.91	1.11	0.96
ITM T3 HE/$\sqrt{10}$HE1	0.50	0.47	0.51	0.49
ITM T4 NH/HE	18.39	18.46	17.81	18.46
ITM T4 $\sqrt{10}$NH1/$\sqrt{10}$HE1	8.80	8.98	8.52	8.83
ITM T4 NH/$\sqrt{10}$NH1	1.12	1.08	1.15	1.10
ITM T4 HE/$\sqrt{10}$HE1	0.54	0.53	0.55	0.53
ITM T5 NH/HE	19.89	19.68	21.11	19.96
ITM T5 $\sqrt{10}$NH1/$\sqrt{10}$HE1	6.53	6.63	6.49	6.57
ITM T5 NH/$\sqrt{10}$NH1	1.05	1.12	1.10	1.09
ITM T5 HE/$\sqrt{10}$HE1	0.34	0.38	0.34	0.36
ITM T6 NH/HE	17.95	18.38	17.77	18.25
ITM T6 $\sqrt{10}$NH1/$\sqrt{10}$HE1	9.12	9.49	8.88	9.25
ITM T6 NH/$\sqrt{10}$NH1	1.12	1.08	1.14	1.10
ITM T6 HE/$\sqrt{10}$HE1	0.57	0.56	0.57	0.56
ITM T7 NH/HE	14.71	15.26	15.01	14.87
ITM T7 $\sqrt{10}$NH1/$\sqrt{10}$HE1	3.51	3.50	3.54	3.47
ITM T7 NH/$\sqrt{10}$NH1	1.49	1.60	1.54	1.54
ITM T7 HE/$\sqrt{10}$HE1	0.36	0.37	0.36	0.36
ITM T8 NH/HE	16.76	18.15	15.44	16.93
ITM T8 $\sqrt{10}$NH1/$\sqrt{10}$HE1	10.08	11.45	9.01	10.60
ITM T8 NH/$\sqrt{10}$NH1	1.06	1.08	1.12	1.08
ITM T8 HE/$\sqrt{10}$HE1	0.64	0.68	0.65	0.67

2.4.3 Time scaling for time-varying correlations

So, the square-root-of-time rule gives reasonable results for unhedged VaR estimates in case of constant correlations, but does not perform well for the hedged option and is highly dependent on the hedge effectiveness over different horizons. For time-varying correlations, the square-root-of-time rule can lead to serious under- and overestimation of the ten-day $\text{VaR}^{0.99}$. Therefore, this section examines whether there is another scaling horizon adequate enough to account for the time-varying correlation term structure such that the under- or overestimation by the square-root-of-time rule remains within 10 percent of the ten-day $\text{VaR}^{0.99}$. The square-root-of-time rule has been examined for the horizon of ten days with the following scalings of $\sqrt{10/j}\times\text{VaR}_{j-day}$ for $j = 1, \ldots, 10$ to see whether any type of scaling is valid. The results show that for correlation term structures with initially positive correlations (T3, T4, T6, and T8), the square-root-of-time rule by scaling the one-day VaR is still reasonable for j of 7 or higher for both hedged and unhedged options. However, the results for the correlation term structures with initially negative correlations (T1, T2, T5 and T7) show that the γ's become close to a value of 1 only for a scaling of j of 9 or more in most cases of the hedged option position. Hence, the scaling of the VaR can only be applied for a time horizon of one to three days in the future for the cases considered here.

2.4.4 Multivariate baskets with time-varying correlations

The simulation study in this chapter considers a basket option with two underlying assets for simplicity, since there is only one correlation term structure involved in the analysis. However, in practice, basket options are often traded on baskets consisting of more than two assets. Longin and Solnik (2001) find that during bear markets, two assets can become increasingly correlated, whereas these assets might have lower correlations in other market conditions. If many assets simultaneously become highly correlated during financial crises, not only do the pairwise correlations have an effect on a basket option on multiple assets, but an interaction effect between different correlation term structures could also exist. According to Fengler and Schwendner (2004), multiple assets in a basket can provide a diversification effect but could also introduce additional risks of unknown correlations, and as correlations between assets rise, the diversification effect will be diminished.

If the purpose of including many assets in a basket is to benefit from a diversification effect, the portfolio of the selected underlying assets should have low correlations in normal market conditions. In that case, the results in this chapter for the correlation term structures with initially low correlations (such as T1, T2) are likely to be more relevant to basket options in practice than the other specifications considered. Moreover, the hedging of the basket options with many assets could be much more difficult. Since the performance of the square-root-of-time rule is highly dependent on the hedge effectiveness, this can have a negative impact on the accuracy of the scaled one-day VaR estimates as an approximation for the ten-day VaR. The basket option does not only react to changes in one correlation term structure as in the case of two assets, but it will also react to changes in $\frac{1}{2}N(N-1)$ different correlation term structures and possibly to interaction effects between these correlation term structures for N assets ($N > 2$). To what extent opposite correlation movements in the basket will offset each other will also depend on the relative importance (i.e., basket weights) of the corresponding assets in the basket. A more detailed analysis of high-dimensional basket options is beyond the scope of this study and is left for further research.

2.5 Conclusions

The purpose of this chapter is to analyze the effect of time-varying correlation term structures on pricing and hedging of basket options as well as on the risk measures VaR and CVaR. First, the benchmark with constant correlations has been used to examine the effect of the sign and size of correlations, and the results are as follows. A surprising result is that basket options on well-diversified portfolio of assets (i.e., negatively correlated assets) are relatively much more sensitive to correlation changes than on a positively correlated portfolio of assets. This implies that sudden changes in the financial markets can lead to large losses for a basket option with well-diversified assets, which cannot be captured by the VaR when using the square-root-of-time rule. The ATM and OTM basket option prices react asymmetrically to positive and negative correlation changes, where a change in negative correlations has a higher impact on the option price than a change in positive correlations of the same magnitude. As a result, the corresponding VaR and CVaR estimates are also more sensitive to changes in negative correlations. The option price of a far ITM option does not differ largely for different correlation values. However, the risk measures of the far ITM basket show that the risk of potential loss increases substantially as the correlation increases. Thus, the risk of potential loss of basket options cannot be entirely captured by the value of the option, but extensive stress testing is needed to reveal these risks.

Second, dynamic delta hedging can reduce the VaR substantially. The specific type of time variation of the correlation term structure is essential to the effectiveness of the dynamic delta hedge of the basket option. Hence, one cannot use the average constant correlation for the computation of the VaR of a hedged basket option with time-varying correlations. Thus, the time of occurrence of a jump during the life of the option is very relevant for the risk of potential loss, even though the basket option payoff is based on the value of the underlying assets at the time of the expiration.

Third, the CVaR can differ from the VaR estimates substantially, especially for the time-varying negative correlations with a jump at the start of the option life and for estimates at the 5 percent significance level. Therefore, providing VaR estimates might not be sufficient, and the coherent measure CVaR can give the additional information needed in certain market conditions.

Finally, the square-root-of-time rule leads to underestimation of the unhedged ten-day VaR in most cases considered, but the underestimation ratio still remains below the regulatory stress factor minimum of 3 for constant correlation term structure. The scaled one-day VaR can deviate heavily from the ten-day VaR for the hedged option when the difference in the hedge effectiveness for the one-day and ten-day horizon is large. The square-root-of-time rule performs relatively well for unhedged ITM basket options. However, applying the square-root-of-time rule can lead to large deviations from the ten-day VaR for (un)hedged ATM and OTM options. In general, the estimates do not improve very much by using different scaling horizons, unless the scaling is applied for square-root-of-time factor close to 1. In these cases, stress testing is much more important for determining an accurate VaR estimate.

Chapter 3

Correlation Bias: Sign and Size Impact on Basket Options

3.1 Introduction

The recent turbulence on the financial markets led to a large number of value-at-risk (VaR) exceedances by many banks in the third quarter of 2007, preceded by a long period of hardly any exceedances. Practitioners ascribe the VaR exceedance clustering to changing market conditions, with sudden increases in asset correlations, and VaR models that were unable to adjust adequately to these changes, see Campbell and Chen (2008). They discuss that some banks discovered that the correlations between asset classes were higher than assumed in their models and correction of these errors had implications for the VaR values. These events clearly demonstrate the need for a closer examination of the impact of sudden correlation increases and errors on risk measures such as VaR.

This study examines the effect of correlation errors for basket options with a time-varying correlation term structure containing a sudden correlation increase. The basket option is a relatively simple option that depends on correlations between multiple assets. There is a wide literature on multivariate correlation models (for a survey see e.g. Bauwens, Laurent and Rombouts (2006)). True asset return correlations are unobservable, hence it is likely that most correlation estimates from covariance models contain estimation errors. Therefore, this chapter takes the following approach to this problem by assuming that estimation errors are present to some extent, regardless of the model used to estimate the correlations. The goal of this chapter is to address the following question: Given a certain size of correlation error, how large is the effect on the risk measures such as (conditional) value-at-risk for basket options for different moneyness, hedging and the use of the square-root-of-time rule?

The approach taken in this chapter is independent of specific correlation models and assumes implicitly that most correlation models are able to produce correlation estimates that are in line with the general pattern of the true time-varying correlations. Wong (2006) has examined the impact of several time-varying correlation term structures on risk measures of a basket option, but does not discuss the impact of correlation errors. This chapter extends this research by quantifying the impact of a correlation bias on the risk measures of a basket option when the correlation term structure contains a jump upward. Practitioners have observed a sudden increase in asset correlations during the turbulent second half of 2007. To my knowledge, the impact of a given size of correlation bias of a time-varying correlation term structure on the risk measures of a basket option for different moneyness levels, hedging and the square-root-of-time rule has not yet been investigated.

Gibson and Boyer (1998) and Byström (2002) also study correlation estimation in relation to option pricing, but evaluate forecast performance of specific multivariate correlation models using options. Skintzi, Skiadopoulos and Refenes (2005) examine correlation errors on European options and use random errors to model the correlation errors. However, these papers do not quantify the size of the VaR misestimation for a certain size of correlation bias and the performance of the square-root-of-time rule in the presence of correlation errors.

The following results of the simulation study of a basket option highlight the risk measure misestimation issues due to correlation errors with regard to option moneyness, hedging and the square-root-of-time rule.

First, the misestimation effect on the risk measures due to correlation errors can be characterized in two ways: The size and asymmetry effect. The size of the risk measure misestimation of the correlation jump specification for the unhedged at-the-money (ATM) and out-of-the-money (OTM) options is larger than for the unhedged in-the-money (ITM) position, implying that unhedged ITM options are relatively less sensitive to correlation errors. The size effect is the largest when delta hedging is applied to the option position, which shows that hedged options are particularly vulnerable to correlation errors. When hedging is applied, a correlation bias not only affects the option price but also the option delta and hence the hedge, thus causing a larger risk measure sensitivity for the hedged position with respect to correlation errors.

Second, an asymmetric misestimation effect occurs for unhedged ATM and OTM options, where a negative correlation bias creates the largest deviations from the unbiased risk measures. Applying hedging increases the asymmetric effect of misestimation for all moneyness levels. For example, a correlation error of -0.1 can already cause an average misestimation of about 26 percent for the VaR at the 99 percent confidence level for the hedged ITM basket option, whereas a correlation bias of +0.1 of the same position causes an average deviation of about 9 percent from the unbiased risk measures.

Finally, the square-root-of-time rule does not adjust properly to changing market parameters and the use of this rule could lead to consistent underestimation of the unbiased ten-day risk measures despite changing signs of the correlation bias for many of the cases considered in this study. Hence, the use of this rule may render very misleading results.

This study proceeds as follows. The recent financial market turmoil in the latter half of 2007 and its implications for the risk measure VaR will be discussed in Section 3.2. The simulation framework to compute the effects of a correlation bias of a time-varying correlation term structure is described in Section 3.3.

3.2 Financial market turbulence

The turmoil in the financial markets that started during the summer of 2007 turned out to be one of the biggest financial crises of the last decades. Many banks admitted to risk management failures and have reported a large number of VaR exceedances in the third quarter of 2007. This clearly demonstrates the need for further research on the current industry standard risk measure VaR. Campbell and Chen (2008) and Campbell (2008) discuss that there have been a large number of VaR exceedances during the third quarter of 2007. Whereas in theory there should have been no more than two or three VaR exceedances in a year, some banks have had more than ten exceedances in the third quarter of 2007 after years of no exceedances at all in relatively tranquil markets. This observation of VaR exceedance clustering — either no exceedances or many in a short time period — shows that VaR is not very adequate for extreme financial markets. Practitioners provide the following reasons for the VaR exceedance clustering in Campbell and Chen (2008): The correlations between assets have unexpectedly risen in a suddenly high volatility environment and the VaR models were not able to adjust properly to reflect the changing market conditions. Some banks discovered that the correlations between assets were higher than assumed in their model and the correction of the correlation errors resulted in changes in VaR.

VaR has been very popular with practitioners and the historical simulation method is often used to compute VaR based on a sample of historical daily returns. Using historical data to estimate the most recent VaR has shortcomings. First, the choice of the sample size of the historical returns is the first problem. A too long sample size can be problematic, because the historical returns may not be able to capture the most recent market movements. Moreover, there may not be enough data available of new financial products. A short sample size could incorporate current events better, but the short historical data sample may not contain extreme events to give an appropriate risk measure for tail events. Moreover, computing VaR using historical data implicitly implies that one only tests for crises that have already occurred and is not forward looking, whereas most financial crises have specific characteristics that have not yet been encountered before. Moreover, practitioners often use the square-root-of-time rule to compute the ten-day VaR by multiplying the one-day VaR with $\sqrt{10}$. This rule can give misleading results for financial derivatives, since scaled empirical one-day returns may not capture the risks that can occur in a ten-day horizon during turbulent markets.

Using Monte Carlo simulation and stress testing to compute the VaR is a method that can overcome the lack of data, the difficult choice of sample size of the historical data and can rapidly be adjusted to reflect current market conditions as well as future market conditions by stress testing. A disadvantage of Monte Carlo simulation is the relatively large computational burden compared to other methods. The two-asset simulation framework discussed in the following section provides results that highlight the risk measure misestimation problems encountered with correlation errors and the use of the square-root-of-time rule.

3.3 Simulation methods

This section uses an approach similar to the framework used in Wong (2006), where a basket option with two underlying assets is used to study the impact of correlation term structures on risk measures. The underlying asset price dynamics are as follows

$$\begin{aligned} dS_{1,t} &= \mu_1 S_{1,t} dt + \sigma_1 S_{1,t} dW_{1,t} \\ dS_{2,t} &= \mu_2 S_{2,t} dt + \sigma_2 S_{2,t} dW_{2,t} \\ dW_{1,t} dW_{2,t} &= \rho_t dt, \end{aligned} \tag{3.1}$$

where $W_{1,t}$ and $W_{2,t}$ are correlated Brownian motions with correlations ρ_t at time t. Correlation jumps are an often occurring empirical feature. When using basket options for diversification purposes, the weighted portfolio of assets should contain underlying assets with low correlations in normal market conditions. The worst fear often mentioned is that investment diversification is least available when most needed. Recent market conditions support this stylized fact: Campbell and Chen (2008) report that during the financial turmoil that began in the third quarter of 2007, the correlations between asset classes has risen unexpectedly rendering diminished diversification effects. Therefore, this chapter will consider the following correlation term structure with a sudden increase in correlations ρ_t given in Figure 3.1 by

$$\begin{aligned} \rho_t = -0.5 \quad &+ \quad \left[0.5\left(t - \frac{1}{21}T\right)252\right]\mathbf{I}_{\left(\frac{1}{21}T < t \leq \frac{3}{21}T\right)} \\ &+ \quad \left[1 - 0.5\left(t - \frac{3}{21}T\right)252\right]\mathbf{I}_{\left(\frac{3}{21}T < t \leq \frac{5}{21}T\right)} \end{aligned} \tag{3.2}$$

The analysis focuses on this specification with a jump in the correlation term structure, since recent financial markets have seen a sudden increase in asset correlations during the turbulent second half of 2007. In practice, the true asset correlations have to be estimated using correlation models and historical data. It is quite likely that most correlation models are able to produce correlation estimates that reflect the general time-varying pattern of the true correlations, but the output of different models can differ in their correlation values. Therefore, the correlation estimates $\hat{\rho}_t$ will be modeled as a perturbation of the entire correlation term structure by the value θ such that

$$\hat{\rho}_t = \rho_t + \theta, \qquad \text{for } \theta \in \{-0.3, -0.2, -0.1, 0, 0.1, 0.2, 0.3\}.$$

If $\theta > 0$ the correlation estimates $\hat{\rho}_t$ overestimate the true correlations ρ_t, if $\theta < 0$ the true correlations are underestimated. The choice of θ values ensures that the perturbed correlation term structure $\hat{\rho}_t$ remains between -1 and +1. The correlation estimates are unbiased if $\theta = 0$ (no misestimation) as depicted in Figure 3.1. Skintzi, Skiadopoulos and Refenes (2005) model misestimation of constant correlations by adding randomized errors to the true correlations. However, it is difficult to analyze the size impact of correlation errors when using randomized errors. The approach in this chapter differs from their approach by perturbing the entire correlation term structure and this has two advantages: it takes correlation estimation errors into account while retaining the overall time-varying correlation pattern, and one can look at the impact of a given specific bias size θ on the risk measures. The basket option in the simulation has a

Figure 3.1: Correlation term structure with estimation error θ

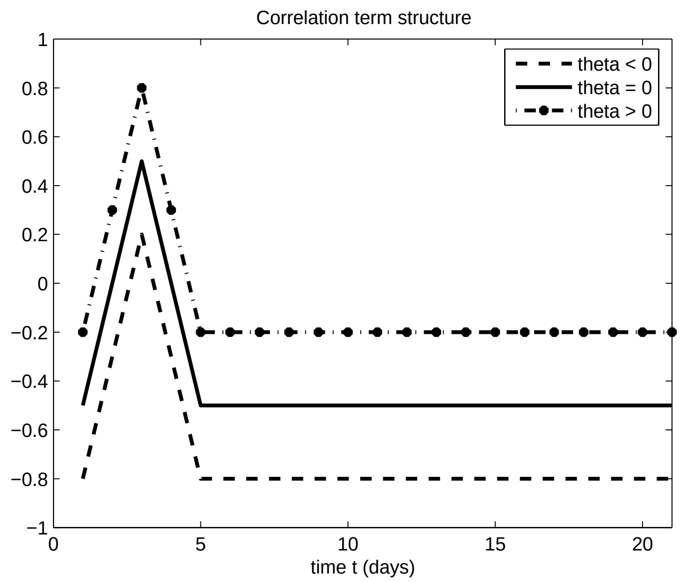

maturity of one month in trading days and the simulation parameters are given in Table 3.1. Let C_T denote the payoff of the basket option at expiry and let $V_t(\hat{S}_i)$ denote the option price at time t as a function of the underlying asset price $\hat{S}_i$ generated using $\hat{\rho}$. When there is no correlation bias (i.e. $\theta = 0$), then $\hat{S}_i$ equals S_i. C_T and V_t are given by

$$C_T = (w_1 \hat{S}_{1,T} + w_2 \hat{S}_{2,T} - K)^+ \tag{3.3}$$

$$V_t = e^{-r(T-t)} \mathbb{E}_{\mathbb{Q}} \left[\frac{1}{M} \sum_{m=1}^{M} C_0^{(m)} \right], \tag{3.4}$$

where M is the number of Monte Carlo simulations and $\mathbb{Q}$ is the risk-neutral measure. The delta of the option can be estimated using the central difference method. The option delta with respect to asset i is given by

$$\Delta_i = \frac{\partial V}{\partial \hat{S}_i} \approx \frac{V(\hat{S}_i + h) - V(\hat{S}_i - h)}{2h}. \tag{3.5}$$

Table 3.1: Option parameters

VaR horizon	10/252	$[\mu_1, \mu_2]$	[0.1,0.1]
Maturity option T	21/252	r	0.05
dt	1/252	K_{ITM}	95
Nr. of assets	2	K_{ATM}	100
$[S_{1,0}, S_{2,0}]$	[100,100]	K_{OTM}	105
$[\sigma_1, \sigma_2]$	[0.35,0.35]	h in (3.5)	0.01
$[\omega_1, \omega_2]$	[1/2,1/2]	M	5000

The risks of the basket option position are assessed using the risk measure value-at risk (VaR), since it is a widely used measure for quantifying potential losses of asset portfolios at a certain confidence level α (conventionally, 95% or 99%). Let X be the simulated profit and loss realizations, then VaR^α is defined by

$$\mathbb{P}(X \leq VaR^\alpha) = 1 - \alpha. \tag{3.6}$$

However, VaR does not give information about the size of the losses once it has been exceeded. Therefore, the conditional value-at-risk (CVaR) will also be computed, because it provides an average value for the tail risks of the option position and is a coherent risk measure (see Artzner et al. 1999) in contrast to VaR. Let X be the profit and loss realizations of the simulations. The CVaR (also called expected shortfall) is defined for a given confidence level α as

$$CVaR^{\alpha} = \mathbb{E}[X|X \leq VaR^{\alpha}]. \tag{3.7}$$

3.4 Results

The impact of correlation bias θ for time-varying correlation term structure containing a jump will be discussed using the results from the simulation experiment. This section discusses the impact of correlation errors on the risk measures considering the size and the sign of θ, option moneyness, hedging and the square-root-of-time rule.

3.4.1 Correlation bias: Size and sign effect

This section presents the results of the simulation experiment, where the impact of a correlation bias on the risk measures is examined by perturbing the correlation term structure over a range of θ values. To quantify the average impact of a marginal correlation bias, Table 3.2 gives the average percentage misestimation of the ten-day risk measures per 0.1 of θ (corresponding to the detailed results given in Tables 3.3 and 3.4).

The results show that the risk measure errors due to biased correlations can be characterized in two ways: Size and asymmetry of the misestimation effect. The size of the misestimation effect of risk measures indicates to what extent the option is sensitive to a correlation bias. The larger the size effect, the less robust the risk measure is with respect to correlation errors. If the misestimation effect is asymmetric in θ, then a positive correlation bias has a different impact on risk measures than a negative bias of the same magnitude, e.g. there is a sign effect of the correlation bias.

The size of the percentage estimation error of the risk measures is increasing in the size of the correlation bias $|\theta|$ as shown in Tables 3.3 and 3.4. Moreover, the size of the risk measure misestimation effect increases substantially for delta-hedged options. The sign of the correlation bias is essential when hedging is applied, since the asymmetric effect of a negative θ is larger than for a positive θ compared to the unhedged case.

Table 3.2: Average (ten-day) risk measure misestimation per θ of 0.1
Given below are the average and percentage deviation from the unbiased risk measures per θ change of 0.1 for θ defined in (3.3). For example, a positive percentage denotes an overestimation of the potential loss indicated by the unbiased risk measures and the average deviation shows the amount by which the potential loss is overestimated.

	Average percentage deviation				*Average deviation*			
No hedging	$VaR^{0.99}$	$VaR^{0.95}$	$CVaR^{0.99}$	$CVaR^{0.95}$	$VaR^{0.99}$	$VaR^{0.95}$	$CVaR^{0.99}$	$CVaR^{0.95}$
ITM $\theta < 0$	1.1	1.5	0.9	1.2	$-1.02\cdot10^4$	$-1.01\cdot10^4$	$-1.03\cdot10^4$	$-1.03\cdot10^4$
ITM $\theta > 0$	−1.1	−1.5	−1.0	−1.3	$1.05\cdot10^4$	$1.02\cdot10^4$	$1.06\cdot10^4$	$1.05\cdot10^4$
ATM $\theta < 0$	2.7	3.9	2.3	3.1	$-2.13\cdot10^4$	$-1.99\cdot10^4$	$-2.11\cdot10^4$	$-2.07\cdot10^4$
ATM $\theta > 0$	−2.0	−2.6	−1.7	−2.2	$1.57\cdot10^4$	$1.34\cdot10^4$	$1.59\cdot10^4$	$1.47\cdot10^4$
OTM $\theta < 0$	1.8	−0.1	1.8	1.3	$-8.69\cdot10^3$	$2.53\cdot10^2$	$-1.10\cdot10^4$	$-5.08\cdot10^3$
OTM $\theta > 0$	−1.3	−0.3	−1.3	−1.1	$6.12\cdot10^3$	$6.27\cdot10^2$	$8.19\cdot10^3$	$4.07\cdot10^3$
Hedging								
ITM $\theta < 0$	26.5	20.4	19.4	20.4	$-1.73\cdot10^4$	$-8.00\cdot10^3$	$-1.86\cdot10^4$	$-1.20\cdot10^4$
ITM $\theta > 0$	−9.2	−13.5	−10.4	−11.7	$6.01\cdot10^3$	$5.30\cdot10^3$	$9.99\cdot10^3$	$6.86\cdot10^3$
ATM $\theta < 0$	18.1	28.4	19.5	22.3	$-1.74\cdot10^4$	$-1.68\cdot10^4$	$-2.18\cdot10^4$	$-1.81\cdot10^4$
ATM $\theta > 0$	−11.2	−14.7	−10.1	−12.7	$1.07\cdot10^4$	$8.72\cdot10^3$	$1.13\cdot10^4$	$1.03\cdot10^4$
OTM $\theta < 0$	22.5	28.2	20.2	24.4	$-1.81\cdot10^4$	$-1.32\cdot10^4$	$-2.01\cdot10^4$	$-1.65\cdot10^4$
OTM $\theta > 0$	−12.3	−16.5	−11.7	−13.4	$9.90\cdot10^3$	$7.72\cdot10^3$	$1.16\cdot10^4$	$9.06\cdot10^3$

Table 3.2 shows that the average impact of a value of 0.1 for θ leads to a deviation from the unbiased risk measures of the unhedged option of approximately 0 to 4 percent for $\theta < 0$ and about 0 to 2 percent for $\theta > 0$, so there is an asymmetric impact of a correlation bias. When the option is delta-hedged, the average impact per 0.1 value of θ is approximately 20 to 28 percent for $\theta < 0$ across different moneyness levels, while the equivalent value for a positive θ deviates about 9 to 16 percent from the unbiased hedged risk measures. Hence, both the size and asymmetric risk measure misestimation effect of the hedged option are much larger than for the unhedged option and this will be discussed in more detail in Section 3.4.3. Across all moneyness levels, a negative correlation bias leads to overestimation of the unbiased risk measures and a positive correlation bias leads to underestimation of the unbiased potential loss. This result is caused by a relatively large change of the option prices over the ten-day period when correlations are lower due to a negative θ. The asymmetric effect is due to the fact that the volatility of the underlying basket changes is more sensitive to changes in low correlations.

This can be seen by looking at the changes of the basket of underlying assets. Following McNeil, Frey and Embrechts (2005) the changes in the underlying basket can be linearized as follows. Let the basket with two underlying assets be $B_t = \sum_{i=1}^{2} w_i S_{i,t}$, and the changes in the underlying basket be $p_{t+1} = B_{t+1} - B_t$. Define the logreturn

$x_{i,t+1} = \ln\left(\frac{S_{i,t+1}}{S_{i,t}}\right)$ for $i = 1, 2$. Then,

$$p_{t+1} = B_{t+1} - B_t = \sum_{i=1}^{2} w_i(S_{i,t+1} - S_{i,t}) = \sum_{i=1}^{2} w_i S_{i,t}(e^{x_{i,t+1}} - 1) \tag{3.8}$$

Linearizing (3.8) gives

$$p^*_{t+1} = \sum_{i=1}^{2} w_i S_{i,t} x_{i,t+1} = B_t \sum_{i=1}^{2} \left(\frac{w_i S_{i,t}}{B_t}\right) x_{i,t+1} = B_t \sum_{i=1}^{2} w^*_i x_{i,t+1} \tag{3.9}$$

Similarly as derived in Appendix 3.A, the logreturn vector $X_t := [x_{1,t}, x_{2,t}]'$ is bivariate normally distributed with covariance matrix Σ_X. Denote $w^* = [w^*_1, w^*_2]'$ and $\tau = (t+1) - t$. Hence, using the properties of the normal distribution, the variance $\sigma^2_{p^*_{t+1}}$ of the linearized profits (3.9) is given in (3.10). For simplicity of discussion, assume that $\rho_u = \rho$ for $u \in [t, t+1]$. Following Taleb (1997) define the correlation vega as the sensitivity of the volatility with respect to the correlation. For p^*_{t+1} it holds that

$$\sigma^2_{p^*_{t+1}} = B_t^2 w^{*'} \Sigma_X w^* = B_t^2 \left[\sum_{i=1}^{2} w^*_i \sigma_i^2 \tau + 2 \sum_{i=1}^{2} \sum_{i<j} w^*_i w^*_j \sigma_i \sigma_j \rho \tau\right]. \tag{3.10}$$

Differentiating (3.10) with respect to ρ gives

$$\frac{\partial \sigma_{p^*_{t+1}}}{\partial \rho} = \frac{1}{2} B_t \left(2 \sum_{i=1}^{2} \sum_{i<j} w^*_i w^*_j \sigma_i \sigma_j \tau\right) \left[\sum_{i=1}^{2} w^*_i \sigma_i^2 \tau + 2 \sum_{i=1}^{2} \sum_{i<j} w^*_i w^*_j \sigma_i \sigma_j \rho \tau\right]^{-\frac{1}{2}} \tag{3.11}$$

$$= \frac{B_t^2 \left(w^*_i w^*_j \sigma_i \sigma_j \tau\right)}{\sigma_{p^*_{t+1}}} \quad , i \neq j \tag{3.12}$$

$$= \frac{w_i w_j S_i S_j \sigma_i \sigma_j \tau}{\sigma_{p^*_{t+1}}} \quad , i \neq j. \tag{3.13}$$

From the correlation vega in (3.11), it is clear that the sensitivity to correlations of the volatility of the linearized profits of the basket declines for a higher correlation ρ, because the numerator increases with a higher ρ. A higher correlation vega when ρ is lower, means that the linearized basket profits will have a greater volatility sensitivity to correlations when correlations are low. This gives rise to an asymmetric effect, where a change in low correlations has a higher effect on the change in the basket than a change of the same size in high correlations.

Finally, the VaR measures have larger misestimation percentages than the CVaR measures in most cases in Table 3.2, hence the correlation errors have a larger percentage impact on the VaR than on the CVaR as shown by the simulation results, since the potential loss indicated by CVaR is relatively larger than VaR. Moreover, VaR gives little information on the size of the potential loss if VaR is exceeded, in contrast to the CVaR that provides an average value for the tail risks. Hence, these reasons indicate that CVaR should be a preferred risk measure to VaR.

3.4.2 Moneyness

The misestimation of the risk measures (C)VaR caused by the correlation bias θ is highly dependent on the option moneyness. Table 3.3 presents the percentage deviation of the ten-day unhedged biased risk measures from the corresponding unbiased values for a range of θ values. The unhedged ATM and OTM option positions have the largest size as well as asymmetric misestimation effect for the risk measures compared to the unhedged ITM option, which responds relatively symmetrically towards a correlation bias (see Table 3.3). For example, a correlation bias ranging from -0.3 to +0.3 causes a misestimation of the unhedged ten-day $\text{VaR}^{0.99}$ varying from about +8.2 percent to -6.0 percent for the ATM option, from +5.5 percent to -3.8 percent for the OTM option and from +3.3 percent to -3.4 percent for the ITM option as given in Table 3.3.

For the ITM option, the incremental misestimation percentage becomes lower when the size of θ increases and equals to around 1 percent. However, for the ATM and OTM option each incremental increase in misestimation becomes slightly larger in case of a negative bias. This implies that the rate at which the misestimation occurs is higher the larger the given negative correlation bias is. This result is more pronounced for the hedged options and will be discussed in the following section.

The risk measures VaR and CVaR react differently towards the same correlation bias. The CVaR measure has relatively lower misestimation percentages than the VaR measures, especially at the 95 percent confidence level and for the ATM option. Hence, the CVaR seems to be more robust towards correlation errors compared to VaR.

Table 3.3: Percentage bias for risk measures of unhedged option

	$\text{VaR}^{0.99}$	$\text{VaR}^{0.95}$	$\text{CVaR}^{0.99}$	$\text{CVaR}^{0.95}$
θ		*In-the-money option*		
−0.3	3.28	4.60	2.84	3.72
−0.2	2.23	3.16	1.94	2.54
−0.1	1.13	1.58	0.98	1.28
0	0.00	0.00	0.00	0.00
0.1	−1.14	−1.57	−0.98	−1.29
0.2	−2.26	−3.10	−1.96	−2.56
0.3	−3.37	−4.61	−2.93	−3.80
θ		*At-the-money option*		
−0.3	8.18	11.72	6.82	9.18
−0.2	5.13	7.18	4.27	5.70
−0.1	2.43	3.34	2.03	2.67
0	0.00	0.00	0.00	0.00
0.1	−2.20	−2.92	−1.85	−2.39
0.2	−4.20	−5.51	−3.56	−4.55
0.3	−6.05	−7.86	−5.14	−6.52
θ		*Out-of-the-money option*		
−0.3	5.50	−0.31	5.40	3.97
−0.2	3.36	0.03	3.42	2.52
−0.1	1.58	0.11	1.62	1.21
0	0.00	0.00	0.00	0.00
0.1	−1.41	−0.19	−1.47	−1.13
0.2	−2.68	−0.47	−2.80	−2.18
0.3	−3.87	−0.77	−4.01	−3.18

3.4.3 Hedging

Option positions are often hedged in practice. Therefore, this section discusses the results when delta hedging is applied to a basket option position given in Table 3.4. The size and asymmetry effect of the risk measure misestimation increase substantially compared to the unhedged results given the same size of correlation estimation error. Delta hedging implies that one sells the underlying asset after a decline in its price and buy the underlying asset after a rise in price at rebalancing times in order to match the delta's of the basket option position. If the underlying asset prices are reasonably stable, delta hedging provides a better hedge compared to highly changing prices that cause higher hedging errors. Depending on the changes in the underlying asset price path and the accuracy of the delta values, delta hedging can lead to losses of the underlying position that may be larger than losses on the option position. Accurate

Table 3.4: Percentage bias for risk measures of hedged option

	$\text{VaR}^{0.99}$	$\text{VaR}^{0.95}$	$\text{CVaR}^{0.99}$	$\text{CVaR}^{0.95}$
θ		*In-the-money option*		
−0.3	79.62	61.26	58.32	61.07
−0.2	39.86	36.04	33.78	36.23
−0.1	17.10	17.18	14.17	16.33
0	0.00	0.00	0.00	0.00
0.1	−9.92	−14.53	−11.30	−13.35
0.2	−19.36	−30.24	−22.09	−24.84
0.3	−27.70	−40.59	−31.30	−35.03
θ		*At-the-money option*		
−0.3	54.31	85.06	58.43	66.82
−0.2	26.48	50.66	32.43	37.97
−0.1	12.90	23.51	13.93	17.02
0	0.00	0.00	0.00	0.00
0.1	−10.69	−16.89	−11.73	−14.70
0.2	−21.72	−31.16	−22.02	−27.25
0.3	−33.55	−44.19	−30.36	−38.23
θ		*Out-of-the-money option*		
−0.3	67.35	84.64	60.53	73.12
−0.2	37.92	48.81	35.42	42.77
−0.1	15.45	20.76	16.31	19.38
0	0.00	0.00	0.00	0.00
0.1	−14.41	−19.63	−13.10	−15.22
0.2	−25.86	−33.46	−25.24	−28.58
0.3	−36.90	−49.37	−34.95	−40.08

delta values are essential for delta hedging, which causes delta-hedged positions to be relatively more sensitive to correlation errors. The asymmetric effect of a negative correlation bias is much larger for the hedged position, which implies that underestimation of the correlation term structure causes a much larger misestimation for the hedged option. So, the hedged option positions are most vulnerable to a negative correlation bias when asset correlations suddenly increase, as recently observed in the turbulent financial markets. For example, the risk measure misestimation error caused by $\theta < 0$ is nearly twice as large as for $\theta > 0$ for the ATM option moneyness. Also, the incremental increase in percentage misestimation is becoming much larger for each increase in θ when θ is negative, especially in case of ITM options.

Although, delta hedging has considerable impact on the reduction of the VaR across all moneyness levels, it also leads to results that are more prone to correlation errors. In the unhedged case, a correlation bias only affects option prices. However,

Table 3.5: Ratio of unbiased ten-day risk measures without hedging and biased risk measures with hedging

	$\mathrm{VaR}^{0.99}$	$\mathrm{VaR}^{0.95}$	$\mathrm{CVaR}^{0.99}$	$\mathrm{CVaR}^{0.95}$
θ		*In-the-money option*		
−0.3	8.02	10.47	7.15	8.78
−0.2	10.30	12.41	8.47	10.38
−0.1	12.30	14.41	9.92	12.16
0	14.40	16.89	11.32	14.14
0.1	15.99	19.76	12.77	16.32
0.2	17.86	24.21	14.54	18.81
0.3	19.92	28.42	16.48	21.77
θ		*At-the-money option*		
−0.3	5.28	4.65	5.25	4.99
−0.2	6.44	5.71	6.28	6.04
−0.1	7.21	6.97	7.29	7.12
0	8.14	8.61	8.31	8.33
0.1	9.12	10.36	9.42	9.76
0.2	10.40	12.50	10.66	11.45
0.3	12.25	15.42	11.93	13.48
θ		*Out-of-the-money option*		
−0.3	3.52	2.80	3.83	3.27
−0.2	4.27	3.47	4.54	3.97
−0.1	5.10	4.28	5.28	4.74
0	5.89	5.17	6.14	5.66
0.1	6.88	6.43	7.07	6.68
0.2	7.94	7.76	8.22	7.93
0.3	9.33	10.20	9.44	9.45

Table 3.4 shows that correlation errors seriously deteriorate the accuracy of the hedged risk measures, because a correlation bias not only affects the option prices, but also the delta's of the options. Hedging the position using unbiased parameters reduces risk, since the hedged position is delta-neutral at rebalancing time. As mentioned earlier in this section, biased option prices and thus biased delta's (see (3.5)) can lead to changes in the hedged position that deviate greatly from the unbiased results for highly changing underlying asset prices (relevant for the quantiles of the loss distribution). These errors are reflected in a bias of the risk measures. Hence, accurate correlations are highly important for hedged option positions. Table 3.5 shows the ratios of the unbiased unhedged risk measures to the hedged risk measures with a correlation bias. One can see that despite a large correlation error, delta hedging still provides a considerable risk reduction compared to the no-bias unhedged risk measures. All ratios are

well above 1 and delta hedging is especially beneficial for the ITM option.

3.4.4 Square-root-of-time rule

Campbell and Chen (2008) discuss the observation of the recent clustering of VaR exceedances, where many banks reported VaR exceedances in a short time period, after some years with a relatively low number of VaR exceedances. This VaR clustering occurred in turbulent market conditions with unexpected increases in correlations across assets during the second half of 2007. As mentioned in their article, clustering of VaR exceedances could be caused by VaR models that are not adjusting well to sudden changes in market conditions. This section shows that the use of the square-root-of-time rule may lead to risk measures failing to adjust appropriately to changing market conditions, especially in the presence of a correlation bias θ. Thus, using the square-root-of-time rule could be an explanation for the VaR exceedance clustering seen in the recent financial markets.

According to the square-root-of-time rule the ten-day VaR can be computed by multiplying the one-day value-at-risk with $\sqrt{10}$. However, the theory states that this result only holds under the assumption of normality and for linear exposures. Hence, the risk measure based on scaling of one-day risks may fail to capture the risks of a ten-day horizon in turbulent markets for nonlinear exposures. Yet, the square-root-of-time rule is often used in practice for convenience reasons. Therefore, this section discusses the performance of this rule when (biased) correlations suddenly jump upward as recently observed in the financial markets. There are two potential sources of misestimation in this context: One source stems from the use of the square-root-of-time rule and the second arises when the market parameters used by the practitioner inserted in the models contain a correlation bias. The following ratios are defined for the unhedged as well as hedged option position to measure the impact of correlation errors when using the square-root-of-time rule

$$\gamma_1^j(\theta) = \frac{\text{ten-day VaR}^\alpha(\theta = 0)}{\sqrt{10} \cdot \text{one-day VaR}^\alpha(\theta)} \quad \text{for j} = \{\text{NH, HE}\}, \tag{3.14}$$

$$\gamma_2^j(\theta) = \frac{\text{ten-day VaR}^\alpha(\theta)}{\sqrt{10} \cdot \text{one-day VaR}^\alpha(\theta)} \tag{3.15}$$

where a subscript for γ of 'NH' will denote the unhedged risk measures and subscript 'HE' represents the risk measures for the hedged option position in the following discussion. The ratio γ_1 shows the true performance of the square-root-of-time rule with respect to the unbiased ten-day value-at-risk estimates. The results for γ_1^{NH} and γ_1^{HE}, γ_2^{NH}, γ_2^{HE} are respectively given in Table 3.6 to 3.9. The ratio γ_2 measures the under- or overestimation of the ten-day risk measures by the square-root-of-time rule estimates, where both risk measures are calculated subject to the same correlation bias value θ. In practice one may not know the size of θ in the correlation estimates, γ_2 measures the performance of the square-root-of-time rule under the assumption of the presence of a latent correlation bias.

First, the results in Tables 3.6 and 3.7 for $\gamma_1(\theta = 0)$ show that the square-root-

Table 3.6: Ratio γ_1^{NH} of unbiased ten-day risk measures and biased square-root-of-time rule risk measures (unhedged option)

	$\text{VaR}^{0.99}$	$\text{VaR}^{0.95}$	$\text{CVaR}^{0.99}$	$\text{CVaR}^{0.95}$
θ		*In-the-money option*		
−0.3	1.17	1.26	1.20	1.22
−0.2	1.20	1.30	1.23	1.25
−0.1	1.23	1.33	1.26	1.28
0	1.26	1.37	1.29	1.31
0.1	1.29	1.40	1.31	1.34
0.2	1.31	1.42	1.34	1.37
0.3	1.34	1.45	1.36	1.39
θ		*At-the-money option*		
−0.3	1.41	1.45	1.45	1.45
−0.2	1.46	1.50	1.50	1.49
−0.1	1.50	1.53	1.54	1.53
0	1.53	1.56	1.57	1.56
0.1	1.56	1.58	1.59	1.58
0.2	1.58	1.59	1.61	1.60
0.3	1.59	1.60	1.63	1.62
θ		*Out-of-the-money option*		
−0.3	2.74	2.47	2.87	2.68
−0.2	2.43	2.09	2.55	2.33
−0.1	2.22	1.87	2.36	2.12
0	2.07	1.74	2.23	1.99
0.1	1.96	1.64	2.14	1.89
0.2	1.88	1.56	2.06	1.82
0.3	1.81	1.50	2.01	1.76

of-time rule highly underestimates the ten-day risk measures for the unhedged and even more for the hedged option position in case of zero correlation bias. This result is consistent with the discussion in Wong (2006). For the ITM option this underestimation ranges from 26 percent to 37 percent in case of no hedging and from 38 to 68 percent for the hedged position across the four (C)VaR risk measures. Moreover, this underestimation is even worse for the unhedged ATM and OTM option position, respectively more than 50 percent and more than a factor of 2. Considering the relatively short option maturity in this simulation experiment, this underestimation of risk over a ten-day horizon by the square-root-of-time rule is quite large.

Second, the misestimation of risk measures caused by the use of the square-root-of-

Table 3.7: Ratio γ_1^{HE} of unbiased ten-day risk measures and biased square-root-of-time rule risk measures (hedged option)

	$VaR^{0.99}$	$VaR^{0.95}$	$CVaR^{0.99}$	$CVaR^{0.95}$
θ		*In-the-money option*		
−0.3	1.43	1.49	1.66	1.51
−0.2	1.39	1.44	1.65	1.47
−0.1	1.38	1.44	1.66	1.47
0	1.38	1.50	1.68	1.49
0.1	1.42	1.49	1.72	1.52
0.2	1.48	1.51	1.75	1.56
0.3	1.50	1.55	1.79	1.60
θ		*At-the-money option*		
−0.3	0.82	0.89	0.81	0.84
−0.2	0.94	1.05	0.94	0.98
−0.1	1.08	1.22	1.06	1.12
0	1.25	1.37	1.17	1.24
0.1	1.37	1.51	1.28	1.35
0.2	1.43	1.61	1.37	1.45
0.3	1.50	1.67	1.45	1.55
θ		*Out-of-the-money option*		
−0.3	1.26	1.37	1.12	1.25
−0.2	1.29	1.37	1.21	1.28
−0.1	1.32	1.41	1.30	1.33
0	1.38	1.47	1.37	1.41
0.1	1.44	1.53	1.45	1.48
0.2	1.49	1.63	1.51	1.55
0.3	1.55	1.70	1.57	1.61

time rule is more severe than the effect of a correlation bias for ten-day VaR estimates as shown in Tables 3.6 and 3.7. Using this rule results in consistent underestimation of the unbiased ten-day risk measures regardless of the sign of θ in nearly all cases considered here (with the exception of the hedged ATM option). For example, Table 3.6 shows that the square-root-of-time rule underestimates the unhedged $\text{VaR}^{0.99}$ of the ITM position by around 17 percent to 34 percent, by 40 to 60 percent for the ATM position and even by a factor of 2.7 to 1.8 for the OTM option when varying θ from -0.3 to +0.3. Basket option prices and delta's are under- or overestimated depending on the sign of θ and hence the misestimation of biased ten-day risk measures should also contain a sign change (see Tables 3.3 and 3.4). However, the consistent underestimation of the unbiased risk measures by the square-root-of-time rule despite changing correlation estimates $\hat{\rho}$ implies that using this rule in practice may lead to a clustering of VaR exceedances over time in turbulent markets with correlation suddenly changing over time. The square-root-of-time rule leads to consistent underestimation if the correlations increase within the ten-day VaR horizon after the first day, since higher correlations of the underlying assets leads to higher volatility and risk of the option position. This can be seen as follows.

Define the logreturns of the underlying assets as $y_{i,t} := \ln\left(\frac{S_{i,t}}{S_{1,0}}\right)$ for $i = 1, 2$ and ρ_t is a deterministic function of time $t \in [0, T]$. The logreturns are bivariate normally distributed such that

$$Y_t := \begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} \quad \sim \quad \mathcal{N}\left(\mu_{Y_t}, \Sigma_{Y_t}\right), \quad \text{where}$$

$$\mu_{Y_t} = \begin{bmatrix} (\mu_1 - \frac{1}{2}\sigma_1^2)t \\ (\mu_2 - \frac{1}{2}\sigma_2^2)t \end{bmatrix} \quad \text{and} \quad \Sigma_{Y_t} = \begin{bmatrix} \sigma_1^2 t & \sigma_1\sigma_2 \int_0^t \rho_s ds \\ \sigma_1\sigma_2 \int_0^t \rho_s ds & \sigma_2^2 t \end{bmatrix} \tag{3.16}$$

For details see Appendix 3.A. If correlations are constant over time (i.e. $\rho_t = \rho$), it holds that $\int_0^t \rho_s ds = \rho t$ and $\Sigma_{Y_t} = At$, where $A(i,j) = \{\sigma_i\sigma_j\rho_{ij}\}_{i=1,2}$. Then, the square-root-of-time rule can be applied to the one-day return on the basket of assets to approximate the ten-day return, since $\sqrt{t}Y_1$ has the same covariance matrix as Y_t. Indeed, the simulation results in Wong (2006) confirm that the square-root-of-time rule performs relatively well for constant correlations.

However, it holds that $\int_0^t \rho_s ds > \rho t$ for the time-varying correlations ρ_t in the set $\Omega := \{\rho_t : \rho_t \geq \rho \ \forall t\ ,\ \exists t > 0 \text{ s.t. } \rho_{t+1} > \rho, \text{ for } \rho_t, \rho \in [-1, 1]\}$, which is the case for the correlation term structure considered in this chapter. Hence, the basket volatility of the logreturns is greater than in case of constant correlations, since $\Sigma_{Y_t}(i,j) > A(i,j)t$

for $i \neq j$. Thus, the square-root-of-time rule will consistently underestimate the true ten-day VaR for substantially increasing correlations over time.

Table 3.8: Ratio γ_2^{NH} of biased ten-day risk measures and biased square-root-of-time rule risk measures (unhedged option)

	$VaR^{0.99}$	$VaR^{0.95}$	$CVaR^{0.99}$	$CVaR^{0.95}$
θ		*In-the-money option*		
−0.3	1.20	1.32	1.23	1.26
−0.2	1.23	1.34	1.25	1.28
−0.1	1.25	1.35	1.27	1.30
0	1.26	1.37	1.29	1.31
0.1	1.27	1.38	1.30	1.32
0.2	1.28	1.38	1.31	1.33
0.3	1.29	1.38	1.32	1.34
θ		*At-the-money option*		
−0.3	1.53	1.63	1.55	1.58
−0.2	1.54	1.61	1.56	1.58
−0.1	1.54	1.58	1.57	1.57
0	1.53	1.56	1.57	1.56
0.1	1.52	1.54	1.56	1.54
0.2	1.51	1.51	1.55	1.53
0.3	1.50	1.48	1.55	1.51
θ		*Out-of-the-money option*		
−0.3	2.90	2.46	3.02	2.78
−0.2	2.51	2.09	2.63	2.38
−0.1	2.26	1.87	2.40	2.15
0	2.07	1.74	2.23	1.99
0.1	1.93	1.64	2.10	1.87
0.2	1.83	1.55	2.01	1.78
0.3	1.74	1.49	1.93	1.71

Third, the square-root-of-time rule results can be misleading in some cases when the correlation bias is latent as is usually the case in practice. Consider the case where the practitioner does not know that there is a large positive correlation bias in his estimates for the correlation term structure. He uses his biased correlations to test the performance of the square-root-of-time rule and the results for ratio γ_2 are given in Tables 3.8 and 3.9. For example, a high positive correlation bias of θ of 0.3 gives γ_2^{HE} ratios of nearly 1 in Table 3.9 indicating a good approximation of the hedged ten-day (C)VaR for ATM options by the square-root-of-time rule. However, Table 3.7 shows that the estimate by the square-root-of-time rule highly underestimates the corresponding unbiased risk measures by around 50 percent for the ATM option with

θ equals 0.3. Hence, this example already shows how the square-root-of-time rule can provide misleadingly well results for the biased ten-day risk measures, whereas very large differences occur when compared to the unbiased risk measures.

Table 3.9: Ratio γ_2^{HE} of biased ten-day risk measures and biased square-root-of-time rule risk measures (hedged option)

	$\text{VaR}^{0.99}$	$\text{VaR}^{0.95}$	$\text{CVaR}^{0.99}$	$\text{CVaR}^{0.95}$
θ		*In-the-money option*		
−0.3	2.56	2.41	2.63	2.44
−0.2	1.95	1.96	2.21	2.00
−0.1	1.62	1.69	1.90	1.71
0	1.38	1.50	1.68	1.49
0.1	1.28	1.27	1.52	1.32
0.2	1.19	1.05	1.36	1.18
0.3	1.09	0.92	1.23	1.04
θ		*At-the-money option*		
−0.3	1.26	1.64	1.28	1.40
−0.2	1.19	1.58	1.24	1.36
−0.1	1.22	1.51	1.21	1.31
0	1.25	1.37	1.17	1.24
0.1	1.23	1.26	1.13	1.15
0.2	1.12	1.11	1.07	1.06
0.3	1.00	0.93	1.01	0.96
θ		*Out-of-the-money option*		
−0.3	2.11	2.53	1.80	2.16
−0.2	1.78	2.04	1.64	1.83
−0.1	1.52	1.70	1.51	1.59
0	1.38	1.47	1.37	1.41
0.1	1.24	1.23	1.26	1.25
0.2	1.10	1.08	1.13	1.11
0.3	0.98	0.86	1.02	0.97

3.5 Conclusions

The impact of correlation estimation errors on risk measures have been studied in this chapter, since recent market conditions of sudden correlation increases have led to VaR exceptions clustering. The following results have been found.

First, biased correlations cause a misestimation effect of risk measures with two characteristics: Size and asymmetry effect. These effects depend on the sign of the correlation bias, the moneyness of the option position and whether or not hedging is applied. The larger the size effect, the more sensitive the risk measure is towards a correlation error. This study shows that the VaR is more sensitive to correlation errors than CVaR. Moreover, VaR gives little information on the size of the potential loss if VaR is exceeded, as opposed to the CVaR that provides an average value for the tail risks. Hence, these reasons indicate that CVaR should be preferred to VaR in providing a robust risk measure. The sign of the correlation bias θ is the most important for ATM and OTM options and for hedged positions, since a negative correlation bias causes a higher misestimation effect.

Second, the results show that given a θ of size 0.1 the average impact on the risk measures for the unhedged position is around 0 to 4 percent, whereas the substantially larger average impact for the hedged option is 10 to 28 percent. Hence, hedged option positions are more sensitive correlation errors than unhedged positions, because the correlation errors have a double impact on the position through the option price as well as the hedge parameters. Hedging option positions in the presence of a correlation bias can lead to a large risk measure misestimation. However, the results show that hedging is still worthwhile for the cases considered here, because the risk reduction benefits outweigh the misestimation effects of a correlation bias.

Finally, the use of square-root-of-time rule may provide very misleading risk measures, because the rule does not adjust well to correlation parameter changes. For correlations increasing over time, in most cases considered here the use of this rule leads to consistent underestimation of the risk measure due to a greater basket volatility than is implied by application of the square-root-of-time rule.

3.A Appendix

In this section the joint distribution of the logreturns of the underlying asset process used in Section 3.4.4 is derived. Let ρ_t be a deterministic function of time and $t \in [0, T]$. The underlying asset process in (3.1) is given by

$$\begin{aligned} dS_{1,t} &= \mu_1 S_{1,t} dt + \sigma_1 S_{1,t} dW_{1,t} \\ dS_{2,t} &= \mu_2 S_{2,t} dt + \sigma_2 S_{2,t} dW_{2,t} \\ dW_{1,t} dW_{2,t} &= \rho_t dt \end{aligned}$$

Define $y_{i,t} := \ln\left(\frac{S_{i,t}}{S_{i,0}}\right)$ and $Y_t = [y_{1,t}, y_{2,t}]'$. The joint distribution of the logreturns of the underlying assets Y_t is derived below by using (stochastic) calculus and Ito's lemma for $i = 1, 2$.

$$dy_{i,t} = d \ln S_{i,t} = \left(\mu_i - \frac{1}{2}\sigma_i^2\right) dt + \sigma_i dW_{i,t} \quad \Rightarrow \tag{3.17}$$

$$y_{i,t} = \ln\left(\frac{S_{i,t}}{S_{i,0}}\right) = \int_0^t \left(\mu_i - \frac{1}{2}\sigma_i^2\right) ds + \int_0^t \sigma_i dW_{i,s} \tag{3.18}$$

Using the properties of the Brownian motion $W_{i,t}$, it follows that Y_t is bivariate normally distributed with mean vector μ_{Y_t} and covariance matrix Σ_{Y_t} (using Ito isometry), where

$$\begin{aligned} \mu_{Y_t}(i) &= \mathbb{E}[y_{i,t}] = \mathbb{E}\left[\int_0^t \left(\mu_i - \frac{1}{2}\sigma_i^2\right) ds + \int_0^t \sigma_i dW_{i,s}\right] = \left(\mu_i - \frac{1}{2}\sigma_i^2\right) t \\ \Sigma_{Y_t}(i,i) &= \mathbb{E}[y_{i,t} y_{i,t}] - \mathbb{E}[y_{i,t}]^2 \\ &= \mathbb{E}\left[\left(\int_0^t \left(\mu_i - \frac{1}{2}\sigma_i^2\right) ds + \int_0^t \sigma_i dW_{i,s}\right)^2\right] - \mathbb{E}\left[\int_0^t \left(\mu_i - \frac{1}{2}\sigma_i^2\right) ds + \int_0^t \sigma_i dW_{i,s}\right]^2 \\ &= \mathbb{E}\left[\left(\int_0^t \sigma_i dW_{i,s}\right)^2\right] = \mathbb{E}\left[\left(\int_0^t \sigma_i^2 ds\right)\right] \\ &= \sigma_i^2 t \end{aligned}$$

Similarly, the covariance element $\Sigma_{Y_t}(i,j)$ for $i \neq j$ is

$$\Sigma_{Y_t}(i,j) = \mathbb{E}\left[\left(\int_0^t \left(\mu_i - \frac{1}{2}\sigma_i^2\right) ds + \int_0^t \sigma_i dW_{i,s}\right)\left(\int_0^t \left(\mu_j - \frac{1}{2}\sigma_j^2\right) ds + \int_0^t \sigma_j dW_{j,s}\right)\right]$$

$$
\begin{aligned}
&-\mathbb{E}\left[\int_0^t \left(\mu_i - \frac{1}{2}\sigma_i^2\right) ds + \int_0^t \sigma_i dW_{i,s}\right] \mathbb{E}\left[\int_0^t \left(\mu_j - \frac{1}{2}\sigma_j^2\right) ds + \int_0^t \sigma_j dW_{j,s}\right] \\
&\quad = \mathbb{E}\left[\left(\int_0^t \sigma_i dW_{i,s}\right)\left(\int_0^t \sigma_j dW_{j,s}\right)\right] = \mathbb{E}\left[\left(\int_0^t \sigma_i\sigma_j\rho_s ds\right)\right] = \int_0^t \sigma_i\sigma_j\rho_s ds
\end{aligned}
$$

Thus, Y_t is distributed as

$$
Y_t := \begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} \quad \sim \quad \mathcal{N}\left(\mu_{Y_t}, \Sigma_{Y_t}\right), \quad \text{where}
$$

$$
\mu_{Y_t} = \begin{bmatrix} (\mu_1 - \frac{1}{2}\sigma_1^2)t \\ (\mu_2 - \frac{1}{2}\sigma_2^2)t \end{bmatrix} \qquad \text{and} \qquad \Sigma_{Y_t} = \begin{bmatrix} \sigma_1^2 t & \sigma_1\sigma_2 \int_0^t \rho_s ds \\ \sigma_1\sigma_2 \int_0^t \rho_s ds & \sigma_2^2 t \end{bmatrix}
$$

If correlations are constant over time, i.e. $\rho_t = \rho \quad \forall t$, then $\Sigma_{Y_t}(i,j)$ is equal to $\sigma_i\sigma_j\rho t$. For more details, see e.g. Brigo and Mercurio (2006).

Chapter 4

Time-Varying Market Liquidity and Static Option Hedging

4.1 Introduction

The measurement of liquidity risk of financial assets is essential for risk management as discussed by the recent study of BIS (2006). Liquidity is time varying (e.g. Evans and Lyons (2002)), can exhibit commonality (e.g., Chordia, Roll and Subrahmanyam (2000), Hasbrouck and Seppi (2001), Huberman and Halka (2001)) and market liquidity can decrease substantially in down markets (Chordia, Roll and Subrahmanyam (2001)). Hence, liquidity risk can arise from idiosyncratic as well as systematic liquidity shocks as also discussed in BIS (2006).

In the literature, the liquidity of an asset is often measured by the size of the spread between the bid and ask price of the asset (see e.g., Lyons (2001)). However, the bid and ask quotes of over-the-counter (OTC) derivatives are generally not widely available or might be stale due to infrequent trading. So, it is not straightforward how to measure the liquidity of (OTC) options, as also acknowledged by Brenner, Eldor and Hauser (2001). A widely used (OTC) option is the barrier option, which is often traded in foreign exchange markets. Barrier options are options that exist depending on whether the price of the underlying asset hits the specified barrier (see Merton (1973) and for a survey, see e.g. Zhang (1998)). Barrier options are popular, since they are cheaper than standard options, but still provide the payoff of a standard European option for a certain range of the underlying asset price.

The main purpose of this study is to examine the liquidity of barrier options. How to measure the liquidity of barrier options, when option quotes are not available? Does market liquidity have an impact on the liquidity of the barrier options? What is the impact of time-varying market liquidity characteristics, such as widening, tightening and jumps of market spreads, on the liquidity costs of static barrier option hedging?

To address these questions, this paper presents a static hedging approach to measure the liquidity of barrier options that only requires the use of the underlying asset prices in the framework of Çetin, Jarrow, Protter and Warachka (2006). This approach avoids the need for (unavailable) option quotes and only uses bid and ask quotes of foreign exchange rates (which are relatively easy to obtain). Moreover, static hedging of barrier options, using a portfolio of standard options (e.g. Derman, Ergener, Kani (1995), Carr and Chou (1997)), outperforms dynamic hedging of barrier options (see Nalholm and Poulsen (2006a)), since option sensitivities besides the option delta are hedged as well. Therefore, the liquidity costs of the barrier option are measured by the liquidity costs of the static hedge. Accordingly, Brenner, Eldor and Hauser (2001) document that the illiquidity discount of currency options can be related to the liquidity of foreign exchange spot markets and should be a function of replicating costs of the option.

The second contribution of this paper is that systematic liquidity is taken into account by generalizing the supply curve in the Çetin, Jarrow, Protter and Warachka (2006) model. BIS (2006) documents that financial institutions find it challenging to prepare for systematic liquidity risks and the approach in this paper allows for explicit computation of idiosyncratic and systematic liquidity costs.

Finally, this paper contributes to the literature by examining the systematic liquidity impact on the liquidity premium of the static barrier option hedge for different time-varying market liquidity characteristics.

Recently, Chordia, Roll and Subrahmanyam (2000), Huberman and Halka (2001), Hasbrouck and Seppi (2001), Fleming (2003), Pástor and Stambaugh (2003), Chordia, Sarkar, Subrahmanyam (2005), Acharya and Pedersen (2005) have examined systematic liquidity in stock and bond markets. Brenner, Eldor and Hauser (2001), Frey and Patie (2002), and Çetin, Jarrow, Protter and Warachka (2006) (and references therein) have studied option illiquidity. However, these studies do not discuss the impact of time-varying market liquidity on static barrier option hedging. Studies on static barrier option hedging, such as Derman, Ergener and Kani (1994), Carr and Chou (1997), Nalholm and Poulsen (2006a), do not examine the impact of time-varying market liquidity, although Carr and Chou (1997) acknowledge the importance of liquidity for static hedging.

To my knowledge, the impact of time-varying market liquidity characteristics, such as widening, tightening and jumps of market spreads, on the liquidity costs of static barrier option hedging in the framework of Çetin, Jarrow, Protter and Warachka (2006), has not yet been examined in the literature.

The results of this study are relevant for financial institutions that have sold OTC barrier options and need to be aware of the substantial option liquidity risk for currencies that are highly dependent on market liquidity. The main results are as follows. Barrier options are often traded on foreign exchange markets and four currencies with different liquidity characteristics have been examined. First, the results show that there is evidence that the liquidity of the Danish Krone, Swiss Franc, British Pound and Norwegian Krone depend significantly on market liquidity represented by the liquidity of the Euro. The liquidity of the Danish Krone is mainly market driven, whereas the liquidity of the Norwegian Krone has a high idiosyncratic component. Liquidity costs of the static barrier option hedge can already be substantial for illiquid currencies in tranquil market conditions (benchmark), e.g. ranging from 8.9 percent for an ITM barrier option to 21.2 percent for an OTM barrier option for the Norwegian Krone.

Second, the results reveal that time-varying market liquidity has substantial impact on liquidity costs of static barrier option hedges. Widening market spreads, often associated with financial crises, cause the largest increase in average liquidity costs compared to tightening market spreads and jumps in market spreads. The average liquidity cost percentage even increases by a factor of 3 for the Danish Krone, which has the lowest liquidity costs in tranquil market conditions. Hence, to determine liquidity risk, it is not sufficient to only consider recent liquidity of a currency, but its dependency on the market liquidity is essential as well. Jumps in market spreads have a relatively large effect on the static hedge of out-of-the-money (OTM) barrier options, considering the small number of jumps during the life of the option.

Finally, the magnitude of the impact of market liquidity on static hedges can vary largely across different moneyness levels of the barrier option. Overall, the static hedge for OTM barrier options has the highest average liquidity cost percentages and is relatively most affected by widening market spreads, whereas ITM options are more affected by tightening market spreads.

The remainder of this article proceeds as follows. Section 1 explains the model for measuring liquidity costs. The estimation of the model for the foreign exchange market is discussed in Section 2. Section 3 and 4 proceed with the simulation methods and results for the time-varying market liquidity scenarios and Section 5 concludes.

4.2 Measuring option liquidity

Static hedging of barrier options outperforms dynamic delta hedging (see e.g. Nalholm and Poulsen (2006a)). The approach in this study measures liquidity costs of the static hedge portfolio consisting of standard European options. An advantage of this approach is that the liquidity costs of standard European options can be measured using only the underlying asset prices in the framework of Cetin, Jarrow, Protter and Warachka (2006). They have proposed the following way to price standard European options in an extended Black-Scholes economy with illiquidity. Denote $S(t,0)$ as the price of the underlying asset. Let $S(t,x)$ be the stochastic supply curve of the underlying asset at time $t \in [0,T]$ corresponding to the incurred trade x, where a positive x is a buy trade and a negative x is a sell trade. The liquidity cost of a standard European option with maturity T can be computed using the trades in the underlying asset incurred by delta hedging the standard options. Following Cetin, Jarrow, Protter and Warachka (2006), let the filtered probability space $[\Omega, \mathcal{F}, (F_t)_{0\leq t\leq T}, \mathbb{P}]$ be given. Denote a self-financing trading strategy as $(X_t, Y_t : t \in [0,T], s)$, where X_t is the aggregate position in the asset, Y_t is the aggregate position in the money market account at time t, s denotes the time of liquidation of the replicating portfolio and r is the risk-free interest rate. Denote x and y as the corresponding discrete trading strategies. Let X_t belong to the class of discrete trading strategies such that

$$X_t \in \left\{ x_{s_0} \mathrm{I}_{(s_0)} + \sum_{j=1}^{N} x_{s_j} \mathrm{I}_{(s_{j-1}, s_j]} \right\}, \tag{4.1}$$

where s_j are stopping times for $j = 1, \ldots, N$, $x_{s_j} \in F_{s_{j-1}} \forall j$, $s_0 = 0, s_N = T$ and $s_j > s_{j-1} + \delta$ for fixed $\delta > 0$. The discounted liquidity costs for a discrete trading strategy is given by

$$L_T = \sum_{j=0}^{N-1} \mathrm{e}^{-rs_j} (x_{s_{j+1}} - x_{s_j})[S(s_j, x_{s_{j+1}} - x_{s_j}) - S(s_j, 0)] \tag{4.2}$$

See Cetin, Jarrow and Protter (2004) and Cetin, Jarrow, Protter and Warachka (2006) for more details. The next section proceeds with the motivation for the specification of the supply curve $S(t, x)$.

4.3 The supply curve

To examine the impact of market liquidity explicitly, the specified supply curve has to allow for idiosyncratic effects as well as for systematic effects for each currency. In this section the supply curve $S(t, x)$ will be specified according to foreign exchange market characteristics. Subsequently, the estimation results for the supply curve are presented.

4.3.1 The foreign exchange market

In markets with high transparency, market order information is quickly reflected in the asset prices. However, foreign exchange markets are relatively non-transparent, because trades do not need to be immediately disclosed, unlike in equity or bond markets, see Lyons (2001). Lyons (2001) identifies three essential characteristics of foreign exchange markets: a foreign exchange market typically has a very large trading volume, a large part of the market consists of interdealer trades and it has low trade transparency. According to BIS (2005), the traditional foreign exchange markets had an estimated average daily turnover of 1,880 billion US Dollar in April 2004. The supply curve $S(t,x)$ is specified by taking these market characteristics into account. Due to high trading volume and low market transparency in the foreign exchange market, it is unlikely that a single trade from a delta-hedging strategy will have a large impact on the asset price in foreign exchange markets. For hedging purposes, the involved transaction costs cannot be too large, hence trades for delta hedging of options are likely to be relatively small compared to the volume in foreign exchange markets. Hence, the liquidity of the foreign exchange market is very likely affected by common market factors rather than individual delta-hedging trades.

Recently, Chordia, Roll and Subrahmanyam (2000), Huberman and Halka (2001), Hasbrouck and Seppi (2001), Fleming (2003), Pástor and Stambaugh (2003), Chordia, Sarkar and Subrahmanyam (2005), Acharya and Pedersen (2005) have examined commonality in (time-varying) liquidity of equity and bond markets. However, these studies do not consider commonality in liquidity in foreign exchange markets and the impact on static hedge portfolios of barrier options.

Gibson and Mougeot (2004) acknowledge that a difficulty related to aggregate market liquidity is to define a proxy for the state variable for aggregate market liquidity. The choice of market liquidity variable in this study is based upon the findings of the most recent Triennial Central Bank Survey (see BIS (2005)), which documents that the US Dollar was on one side of 89 percent of all foreign exchange trades and the US Dollar/Euro rate had the largest global turnover of 28 percent in April 2004. Therefore, the common market liquidity $M(t,x)$ will be represented by the liquidity of the US Dollar/Euro exchange rate. Moreover, all exchange rates in this paper will have the US Dollar on one side. To my knowledge, there is no study that examines commonality in liquidity in foreign exchange markets using the liquidity of the USD/EUR as the common factor specified in the supply curve of the model of Cetin, Jarrow, Protter and Warachka (2006).

To specify the supply curve, let $S_i(t,x)$ be either the bid or the offer price of the trade x of exchange rate i. Let $S_i(t,0)$ be the exchange mid rate, where the mid rate is defined as the average of the bid and offer rate. Suppose the common market liquidity factor $M(t,x)$ is given by the liquidity of the dominant market currency pair USD/EUR. Denote the bid (or offer) rate of this dominant currency pair as $S_m(t,x)$ and the exchange mid rate as $S_m(t,0)$. The supply curve in this paper is specified as follows[1]

$$S_i(t,x) = S_i(t,0)\left(e^{sign(x)\alpha_i}\cdot M(t,x)^{\beta_i}\right), \qquad (4.3)$$

$$M(t,x) = \frac{S_m(t,x)}{S_m(t,0)}, \qquad (4.4)$$

where the impact on $S_i(t,x)$ is composed of an idiosyncratic effect measured by α_i and a systematic effect from the liquidity of a leading currency pair in the market is measured by β_i. Due to the high volume in the foreign exchange markets and the relative small impact of delta-hedge trades, the supply curve in the foreign exchange market only depends on the sign of the trade. This supply curve specification allows for comparison of parameters α and β across different currencies, since these are estimated based on proportional liquidity of the currencies and the next section proceeds with the estimation results.

4.3.2 Estimation results

The data consist of daily WM/Reuters foreign exchange ask and mid rates of the Danish Krone (DKK), Euro (EUR), British Pound (GBP), Swiss Franc (CHF), and the Norwegian Krone (NOK). These European currencies correspond to about the same geographical region and time zone, so it is very likely that they are affected by a common market factor. The exchange rates are expressed in USD per unit foreign currency, the data are from October 2004 to March 2006 and obtained from Datastream. The common market liquidity for the supply curve, M_t, is computed using the ask and mid rates of the currency pair USD/EUR. The supply curve has the same parameters for buy and sell trades and thus can be estimated by only using the ask and mid rates of the currencies. Suppose now that x is a buy trade, hence $sign(x)$ is equal to 1. Hence, by taking the logarithms on both sides of equation (4.3), the parameters of currency i

[1] Following Cetin, Jarrow, Protter and Warachka (2006) $sign(x)$ is defined as

$$sign(x) = \begin{cases} +1 & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x < 0 \end{cases}$$

can be estimated as follows

$$\ln\left(\frac{S_i(t,x)}{S_i(t,0)}\right) = \alpha_i + \beta_i \ln\left(\frac{S_m(t,x)}{S_m(t,0)}\right) + \epsilon_i(t), \tag{4.5}$$

where i = CHF, DKK, GBP, NOK and $\epsilon_i(t)$ is the error process. The standard errors of the parameters computed are the Newey and West (1987) heteroscedasticity and autocorrelation consistent standard errors. The left hand side variable of (4.5), the log of the ask/mid ratio, can be interpreted as approximately the proportional ask-mid spread percentile, since the ask/mid ratio is usually close to one. The regression results in

Table 4.1: Estimation and parameters
This table presents the supply curve estimation results and the parameters used for the simulation experiment. The t-statistics are given in parentheses. (*) denotes significant at 5% level and (**) denotes significant at 1% level.

	CHF	DKK	GBP	NOK
Regression results				
α ($\times 10^{-4}$)	1.60**	0.27**	0.95**	2.91**
	(13.0)	(6.7)	(7.8)	(14.8)
β	0.94**	0.99**	0.25*	0.77**
	(9.7)	(31.6)	(2.4)	(5.3)
R^2	0.21	0.67	0.03	0.05
Simulation parameters				
$S(0,0)$	0.800	0.167	1.796	0.149
μ	−0.029	−0.017	−0.023	0.013
r_f (%)	0.90	2.38	5.05	2.07
σ (%)	9.80	8.60	8.20	10.10

Table 4.2: Currency liquidity characteristics

	Low	Medium	High
Idiosyncratic α	GBP, DKK	CHF	NOK
Systematic β	GBP	NOK	DKK, CHF

Table 4.1 show that the log of the ask/mid price ratio of the Danish Krone, the Swiss Franc, the British Pound and the Norwegian Krone depend significantly on the log of the ask/mid price ratio of the market currency Euro. The results of the Danish Krone and the Swiss Franc show that the systematic parameter β (respectively, 0.99 and 0.94)

and R^2 (respectively, 0.71 and 0.21) are relatively high for both currencies, so their liquidity is highly dependent on the market liquidity M_t. However, the idiosyncratic component of the liquidity of both currencies differ considerably. The idiosyncratic parameter α of the Danish Krone has a relatively small value of $0.27 \cdot 10^{-4}$. Hence, the liquidity of the Danish Krone is mainly driven by the market liquidity, whereas the liquidity of the Swiss Franc has a quite high idiosyncratic component with an α estimate of $1.60 \cdot 10^{-4}$.

The α for the Norwegian Krone has a very high value of $2.91 \cdot 10^{-4}$, which is more than ten times as large as the α for the Danish Krone. So, the Norwegian Krone is a relatively illiquid currency with very high bid-ask spreads, mainly due to the idiosyncratic effect in this model. Though, the liquidity of the Norwegian Krone is significantly affected by the market liquidity measure as well, with a β estimate of 0.77.

The British Pound is only moderately affected by the market liquidity of the Euro, but the idiosyncratic parameter is not very large either. This result is in line with intuition, since a major currency pair as the USD/GBP with a high global turnover has a very high liquidity and thus relatively low values for α and β.

The four currencies presented here have different levels of idiosyncratic and systematic liquidity parameters (summarized in Table 4.2) and are used to investigate how changes in the market liquidity measure will affect the liquidity costs of barrier options on currencies with different liquidity characteristics in the subsequent simulation experiment.

4.4 Simulation method

Prior to presenting the simulation results, the next section discusses the static hedging methods and simulation techniques of time-varying bid-ask spreads needed for the simulation experiment.

4.4.1 Static hedging

Two important static hedging methods for barrier options using a portfolio of standard European options are the calendar-spread hedge of Derman, Ergener and Kani (1995) and the strike-spread hedge of Carr and Chou (1997). Subsequently, these methods are described for an up-and-out barrier call option following the approach of Nalholm and Poulsen (2006b). The payoff of an up-and-out barrier call option with strike K, barrier B and maturity T equals the payoff of a standard European call option with the same strike and maturity, if the underlying asset price stays below the barrier value B during the life of the option and becomes worthless if the barrier is hit (see Appendix 4.A for the option formula).

Derman, Ergener and Kani (1995) have proposed a static hedging method using a weighted portfolio of standard European option with strikes at K and B and different maturities. Let $C(S(t,0),t|K,T)$ denote the Black-Scholes (1973) call option price at time t, with strike K and expiry date T and this option is first included in the static hedge. Then, the maturity of the barrier option is divided into N_p matching points t_k, where $t_k < T$ for $k = 1, 2, \ldots, N_p$. The weights of the portfolio of standard European call options can be determined by recursively computing the number of options needed to get a static hedge value of zero at the entire barrier. To match the barrier option at time t_{N_p} if the underlying asset hits the barrier, determine the weight γ_{N_p} in the option with strike B and maturity T by solving

$$\gamma_{N_p} C(B, t_{N_p}|B,T) = -C(B, t_{N_p}|K,T). \tag{4.6}$$

Let τ_k be the expiry date of the call option included in the portfolio corresponding to matching point t_k. So, there are call options at strike levels K and B that mature at T. Nalholm and Poulsen (2006b) denote this as the expiry hedge

$$\Lambda(S(t,0),t|K,T) = C(S(t,0),t|K,T) + \gamma_{N_p} C(S(t,0),t|B,T). \tag{4.7}$$

Then, at matching point t_k, a call option is used with expiry τ_k such that $t_k < \tau_k \leq t_{k+1}$. The position in this call option should be chosen such that the static hedge will have zero value at the barrier B at time t_k to match the barrier option value. The weights can be determined for each matchpoint by solving recursively

$$\gamma_k C(B, t_k|B, \tau_k) = -\Lambda(S(t_k, 0), t_k|K,T) - \sum_{u=k+1}^{N_p-1} \gamma_u C(B, t_k|B, \tau_u). \tag{4.8}$$

The strike-spread approach of Carr and Chou (1997) uses a portfolio of options with expiry date T, but with different strikes. For an up-and-out call, these strikes will be set beyond the barrier. As in the calendar-spread approach, t_k are the matching points for $k = 1, \ldots, N_p$. The hedge portfolio consists of N_p call options with strikes κ_l (for $l = 1, \ldots, N_p$) and corresponding portfolio weight γ_l. To match the option payoff if the barrier is never hit, a standard European option struck at K with maturity T is included. The weights vector $\gamma = (\gamma_1, ..., \gamma_{N_p})'$ for the call options with different strike levels can be determined by solving the system

$$A\gamma = \Pi, \tag{4.9}$$

where $\Pi_k = -C(B, t_k|K, T)$ and $A_{k,l} = C(B, t_k|\kappa_l, T)$. For more details, see Derman, Ergener and Kani (1995), Carr and Chou (1997) and Nalholm and Poulsen (2006b).

4.4.2 Time-varying market spreads

It is widely known that bid-ask spreads of currencies are varying over time and exhibit market characteristics such as discreteness and clustering of spread values (see, e.g., Hasbrouck (1999) and references therein). The implicit tick size of bid and ask quotes is an i.i.d. random variable in the Hasbrouck (1999) model. The bid-ask spreads of the USD/EUR exchange rate in the data sample in Section 2.2 mainly cluster around the values $3 \cdot 10^{-4}$ and $4 \cdot 10^{-4}$. Therefore, the market spreads in the simulation are modeled in a similar way as the implicit tick size in Hasbrouck (1999). The bid-ask spread of the USD/EUR exchange rate at time t, denoted by $BA_M(t)$, is modeled as an i.i.d. random variable in the simulation experiment as follows

$$\text{Case 1 Benchmark:} \quad BA_m^1(t) = \begin{cases} 3 \cdot 10^{-4} & \text{with probability } 0.80 \\ 4 \cdot 10^{-4} & \text{with probability } 0.20, \end{cases} \tag{4.10}$$

Studies such as Chordia, Roll and Subramahnyam (2001) have found that liquidity plummets in down markets. A decline in liquidity can be represented by soaring bid-ask spreads. The minimum and maximum bid-ask spread observed in foreign exchange data often differ by a factor of five or more, see Bessembinder (1994) and Hasbrouck

(1999). In the following simulation experiment, the impact of time-varying market bid-ask spreads on the liquidity cost of an up-and-out barrier call option is examined by simulating scenarios of widening bid-ask spreads (Case 2), tightening bid-ask spreads (Case 3) and jumps (Case 4) using Case 1 in (4.10) as a benchmark. The case of bid-ask spreads widening over time can be interpreted as a pre-crisis market condition and will be simulated by multiplying (4.10) with a factor that is linearly increasing over time from one to five. Tightening bid-ask spreads are analogously simulated with a factor linearly decreasing over time from five to one, and correspond to a possible post-crisis market condition. Finally, Case 4 contains jumps that have a high value and a low probability of occurring and this bid-ask spread scenario is given by

$$\text{Case 4 Jumps in spreads:} \quad BA_m^4(t) = \begin{cases} 3 \cdot 10^{-4} & \text{with probability } 0.80 \\ 4 \cdot 10^{-4} & \text{with probability } 0.15 \\ 20 \cdot 10^{-4} & \text{with probability } 0.05, \end{cases} \tag{4.11}$$

The foreign exchange mid rates $S_i(t,0)$ and the mid rate of the market liquidity factor given by the Euro, $S_m(t,0)$ are simulated by the following processes

$$dS_m(t,0) = \mu_m S_m(t,0)dt + \sigma_m S_m(t,0)dW_m(t) \tag{4.12}$$

$$dS_i(t,0) = \mu_i S_i(t,0)dt + \sigma_i S_i(t,0)dW_i(t), \tag{4.13}$$

where $i =$ CHF, DKK, GBP and NOK, m corresponds to the market liquidity factor, μ and σ are, respectively, the return and volatility of the currency and $W_i(t), W_m(t)$ are independent Brownian motions. Subsequently, the market liquidity factor of (4.4) for each bid-ask spread case is simulated by

$$M(t,x) = \frac{S_m(t,0) + sign(x) \cdot \frac{1}{2} BA_m(t)}{S_m(t,0)}, \tag{4.14}$$

which can be interpreted as the proportional liquidity corresponding to either a buy or a sell (a signed trade) and the numerator is $S_m(t,x)$, which is the mid rate adjusted for either a bid or a ask spread for the signed trade. The bid-ask spreads of currency i will be endogenously determined by the supply curve given in (4.3) using the parameters given in Table 4.1. The parameters μ and σ have been estimated from the data sample in Section 2.2. The maturity of the barrier option is assumed to be nine months in trading days and the barrier value is struck at 15 percent above the initial exchange rate S(0,0). The barrier option is examined for moneyness levels of 5 percent in-the-money (ITM), at-the-money (ATM) and 5 percent out-of-the-money (OTM). The other simulation parameters are based on data of October 1st, 2004 that can be obtained from Datastream. The parameters of the initial stock price is based on the exchange mid rate of the currencies and the interest rates for CHF, DKK and GBP are based on the nine-month Libor rate. There is no Libor rate available for the NOK and the nine-month Norway interbank offer rate has been used instead.

4.5 Simulation Results

The simulation results of the different specifications of time-varying market liquidity are presented and their impact on the liquidity costs of the barrier option will be discussed. There are differences between the results for the calendar-spread (DEK) and the strike-spread (CC) when using the liquidity costs as a performance criterion, but these are not very large. Overall, the liquidity costs as a percentage of the static hedge value are higher for the strike-spread hedge. The absolute average liquidity costs of the strike-spread hedge are higher for initially ITM barrier option, but the calendar-spread hedge has higher average liquidity costs for OTM options under most market scenarios. The initial setup costs for the static hedge are higher for the calendar-spread hedge than for the strike-spread. For a comparison of static hedging methods using the hedging error as a performance criterion, see Nalholm and Poulsen (2006a). The subsequent discussion is based on the results for the strike-spread static hedge.

4.5.1 Benchmark

The simulation results for the benchmark Case 1 are given in Table 4.3. The static hedge for a barrier option on the Danish Krone has the lowest percentage average liquidity cost, followed in increasing order by the British Pound, the Swiss Franc and the Norwegian Krone. Moreover, the average liquidity costs are the highest for initially

ITM options and the lowest for OTM options for each currency, but the OTM option has the highest impact of liquidity costs in percentage terms of the initial value of the static hedge. This result is in line with the result that Çetin, Jarrow, Protter and Warachka (2006) find for standard European options. The average liquidity cost percentage for the Danish Krone varies from 2.6 percent for an ITM barrier option to 7.1 percent for an OTM barrier option, and these costs are slightly higher for the British Pound. The highest percentage liquidity costs of the static hedge are for the Norwegian Krone, which vary from 8.9 percent for the ITM option to 21.2 percent for the OTM option. Thus, the liquidity costs of a static barrier option hedge for a relatively illiquid currency as the Norwegian Krone can already be substantial in the tranquil market conditions of this benchmark Case 1, especially for OTM barrier options. Although, the liquidity cost percentages seem quite high, particularly for the Norwegian Krone, their magnitude is in line with results that can be found in the literature. For example, Brenner, Eldor and Hauser (2001) report empirical illiquidity discounts that vary from 17 percent to 27 percent for an illiquid currency option issued by a central bank.

The systematic percentage of the liquidity costs is nearly the same across different

Table 4.3: Benchmark (Case 1)
The calendar-spread static hedge is indicated by DEK and the strike-spread static hedge by CC. For each method, the results are given for the initial value of the static hedge for the up-and-out barrier call option, the average liquidity costs and the average liquidity costs as a percentage of the initial static hedge value. (*) denotes significant at 5% level and (**) denotes significant at 1% level.

Currency	Strike	DEK			CC		
		Value	Liq. costs	Liq. %	Value	Liq. costs	Liq. %
CHF	ITM	$3.99 \cdot 10^3$	190.2	4.8	$3.87 \cdot 10^3$	193.2	5.0
	ATM	$1.86 \cdot 10^3$	128.2	6.9	$1.77 \cdot 10^3$	128.2	7.3
	OTM	$6.09 \cdot 10^2$	61.9	10.2	$5.50 \cdot 10^2$	59.9	10.9
DKK	ITM	$8.25 \cdot 10^2$	20.3	2.5	$8.12 \cdot 10^2$	20.8	2.6
	ATM	$3.61 \cdot 10^2$	14.5	4.0	$3.51 \cdot 10^2$	14.7	4.2
	OTM	$1.09 \cdot 10^2$	7.3	6.7	$1.02 \cdot 10^2$	7.2	7.1
GBP	ITM	$7.15 \cdot 10^3$	155.7	2.2	$7.09 \cdot 10^3$	160.1	2.3
	ATM	$2.81 \cdot 10^3$	116.3	4.1	$2.76 \cdot 10^3$	117.9	4.3
	OTM	$7.42 \cdot 10^2$	56.4	7.6	$7.11 \cdot 10^2$	56.2	7.9
NOK	ITM	$6.81 \cdot 10^2$	57.4	8.4	$6.61 \cdot 10^2$	58.8	8.9
	ATM	$3.10 \cdot 10^2$	38.9	12.6	$2.95 \cdot 10^2$	39.4	13.3
	OTM	$9.95 \cdot 10^1$	19.3	19.4	$8.98 \cdot 10^1$	19.0	21.2

moneyness levels in this benchmark case. Table 4.7 shows that 82.7 percent of the liquidity costs of the Danish Krone can be ascribed to the systematic liquidity component. Thus, the liquidity costs of the static barrier option hedge on the Danish Krone are mainly driven by market liquidity. The liquidity costs of the Swiss Franc has a systematic component of around 43.3 percent, so the liquidity of the Swiss Franc is not only affected by the market liquidity, but more than half of the liquidity costs are incurred due to the idiosyncratic liquidity component. Although the Norwegian Krone has a quite high systematic β parameter of 0.77, the contribution of the market factor towards the liquidity costs is only 25.7 percent, since the idiosyncratic parameter α is very large relative to other currencies. Hence, the impact of systematic liquidity on the average total liquidity costs of the static barrier option hedge depends on the relative size of the market parameter β compared to α. In the benchmark case, the average liquidity costs of the static barrier hedge for the British Pound consist mainly of idiosyncratic risk, with only about a quarter of the liquidity costs coming from the systematic component. However, this impact of market liquidity can increase rapidly for widening market spreads, often associated with financial crises, as discussed subsequently.

4.5.2 Widening market spreads

The results for widening market bid-ask spreads of the dominant market currency, representing decreasing market liquidity over time, are given in Table 4.4. Overall, wide market spreads at the end of the option life cause a substantial increase in the static hedging percentage liquidity costs. This increase is the largest for the Danish Krone and the Swiss Franc, since the liquidity of these currencies are most affected by market liquidity due to a high systematic parameter β. For example, the liquidity cost percentages for the static hedge for the Danish Krone range from 7.6 percent to 23.1 percent, respectively for the ITM and OTM barrier option, which are about three times as large as in the benchmark case.

Similar to the results of the benchmark, the percentage liquidity costs of the static hedge are the highest for the OTM barrier options compared to other moneyness levels. Moreover, the OTM liquidity cost percentages for each currency increase with the highest factor for increasing market spreads compared to the benchmark. For example, the percentage average liquidity costs for the Swiss Franc increase by a factor of 2.2 for the static hedge of the OTM barrier option, whereas this factor is 2.0 for the ITM barrier option. So, decreasing market liquidity has a relatively large impact on the average liquidity costs of the static hedge of the OTM barrier option due to a high systematic impact.

The systematic component of the total liquidity costs has increased for all currencies, when market spreads are increasing over time. More than 94 percent of the average liquidity costs for the Danish Krone come from the systematic liquidity component, which implies that the liquidity costs of the static hedge for the Danish Krone barrier option is very sensitive to changes in the market liquidity. The percentage liquidity costs for decreasing market liquidity are twice as high as those of the British Pound, whereas these percentages were lower for the Danish Krone than for the British Pound in the benchmark case. Thus, widening market spreads can cause a relatively large increase of static hedging liquidity costs for currencies that are very liquid in calm market conditions, but also highly dependent on the market liquidity. The systematic percentage of the total average liquidity costs for the British Pound and the Norwegian Krone doubled to more than 50 percent in the case of widening market spreads,

Table 4.4: Widening market spreads (Case 2)
The calendar-spread static hedge is indicated by DEK and the strike-spread static hedge by CC. For each method, the results are given for the initial value of the static hedge for the up-and-out barrier call option, the average liquidity costs and the average liquidity costs as a percentage of the initial static hedge value. (*) denotes significant at 5% level and (**) denotes significant at 1% level.

Currency	Strike	DEK			CC		
		Value	Liq. costs	Liq. %	Value	Liq. costs	Liq. %
CHF	ITM	$3.99 \cdot 10^3$	379.3	9.5	$3.87 \cdot 10^3$	389.5	10.1
	ATM	$1.86 \cdot 10^3$	266.0	14.3	$1.77 \cdot 10^3$	269.0	15.2
	OTM	$6.09 \cdot 10^2$	133.6	22.0	$5.50 \cdot 10^2$	131.2	23.8
DKK	ITM	$8.25 \cdot 10^2$	59.6	7.2	$8.12 \cdot 10^2$	61.8	7.6
	ATM	$3.61 \cdot 10^2$	44.8	12.4	$3.51 \cdot 10^2$	45.8	13.0
	OTM	$1.09 \cdot 10^2$	23.5	21.6	$1.02 \cdot 10^2$	23.5	23.1
GBP	ITM	$7.15 \cdot 10^3$	248.6	3.5	$7.09 \cdot 10^3$	257.0	3.6
	ATM	$2.81 \cdot 10^3$	189.4	6.7	$2.76 \cdot 10^3$	193.3	7.0
	OTM	$7.42 \cdot 10^2$	94.0	12.7	$7.11 \cdot 10^2$	94.5	13.3
NOK	ITM	$6.81 \cdot 10^2$	92.8	13.6	$6.61 \cdot 10^2$	96.1	14.5
	ATM	$3.10 \cdot 10^2$	64.9	21.0	$2.95 \cdot 10^2$	66.4	22.5
	OTM	$9.95 \cdot 10^1$	33.2	33.4	$8.98 \cdot 10^1$	33.2	37.0

in contrast to the benchmark case, where the liquidity costs of the British Pound and the Norwegian Krone are mainly incurred due to idiosyncratic liquidity risk. The systematic percentage of the total liquidity costs in Table 4.4 is the highest for OTM barrier options across all currencies. Hence, widening spreads can have a larger effect on the liquidity costs of the static hedge for OTM barrier options compared to other moneyness levels.

4.5.3 Tightening market spreads

Increasing market liquidity is represented by a scenario where initially high market spreads are decreasing over time and the results are given in Table 4.5. The percentage liquidity costs of a static hedge are larger than in the benchmark case, but the increase is not as large as in the case of widening market spreads. These results show that wide market bid-ask spreads near the maturity of the barrier option have a larger impact on the liquidity costs for a static hedge than tightening spreads, which could be due to rapidly changing option delta values near maturity. Similar to the results of widening market spreads, the time-varying market liquidity causes the Danish liquidity costs to be larger than those for the British Pound in contrast to the benchmark, since the liquidity of the Danish Krone is highly dependent upon the market liquidity factor. The Norwegian Krone remains the currency with the largest liquidity costs for the static hedge varying from 12.4 percent for an ITM barrier option to 27.2 percent for an OTM barrier option.

The systematic liquidity risk percentage of the Danish Krone has increased with about 10 percent, as shown in Table 4.5. The largest increase in systematic risk percentage can be found for the British Pound and Norwegian Krone, that is by more than 25 percent compared to the benchmark.

Although, the static hedge for an OTM barrier option still has the highest percentage liquidity costs across moneyness levels, the increase in the percentage liquidity costs are larger for the ITM barrier option in the case of decreasing market spreads compared to the benchmark. For example, the results for the Swiss Franc in Table 4.5 show that the total liquidity cost percentages increase by a factor of 1.7 for the static hedge of the ITM barrier option, whereas this factor is 1.5 for the OTM barrier option relative to the benchmark percentages. Moreover, the systematic percentage of

Table 4.5: Tightening market spreads (Case 3)
The calendar-spread static hedge is indicated by DEK and the strike-spread static hedge by CC. For each method, the results are given for the initial value of the static hedge for the up-and-out barrier call option, the average liquidity costs and the average liquidity costs as a percentage of the initial static hedge value. (*) denotes significant at 5% level and (**) denotes significant at 1% level.

Currency	Strike	DEK			CC		
		Value	Liq. costs	Liq. %	Value	Liq. costs	Liq. %
CHF	ITM	$3.99 \cdot 10^3$	330.7	8.3	$3.87 \cdot 10^3$	331.9	8.6
	ATM	$1.86 \cdot 10^3$	212.8	11.5	$1.77 \cdot 10^3$	209.9	11.9
	OTM	$6.09 \cdot 10^2$	97.7	16.0	$5.50 \cdot 10^2$	92.5	16.8
DKK	ITM	$8.25 \cdot 10^2$	48.1	5.8	$8.12 \cdot 10^2$	48.7	6.0
	ATM	$3.61 \cdot 10^2$	32.4	9.0	$3.51 \cdot 10^2$	32.2	9.2
	OTM	$1.09 \cdot 10^2$	15.3	14.1	$1.02 \cdot 10^2$	14.7	14.4
GBP	ITM	$7.15 \cdot 10^3$	221.7	3.1	$7.09 \cdot 10^3$	226.4	3.2
	ATM	$2.81 \cdot 10^3$	161.9	5.8	$2.76 \cdot 10^3$	162.7	5.9
	OTM	$7.42 \cdot 10^2$	76.5	10.3	$7.11 \cdot 10^2$	75.3	10.6
NOK	ITM	$6.81 \cdot 10^2$	80.7	11.8	$6.61 \cdot 10^2$	81.8	12.4
	ATM	$3.10 \cdot 10^2$	52.8	17.1	$2.95 \cdot 10^2$	52.8	17.9
	OTM	$9.95 \cdot 10^1$	25.2	25.3	$8.98 \cdot 10^1$	24.4	27.2

the total liquidity costs is higher for ITM options than for OTM options when market spreads are tightening over time, and the opposite holds for the case of market spreads widening over time as shown in the previous section. This result is likely due to the large part of the static hedging liquidity costs that is incurred at the initial setup of the hedge for ITM options, whereas the initial setup costs are much lower for OTM options. Therefore, wide spreads at the end of the option life could have a relatively higher systematic impact on the OTM option than for other moneyness levels.

4.5.4 Jumps in market spreads

The results for the case where market liquidity drops on a small number of days and then restores again are given in Table 4.6. The percentage average liquidity costs of the British Pound and the Norwegian Krone increase only slightly, while the Danish Krone has an increase by a factor of 1.2 compared to the benchmark liquidity cost percentages. Considering the fact that the large jumps in bid-ask spreads only occur in 5 percent of the days during the nine-month maturity period of the barrier option, this increase in liquidity costs is quite large. Unlike in Case 2 and 3, the increase is about the same across different moneyness levels.

The increase in the systematic liquidity percentage of the total liquidity costs range from around 3 percent for the Danish Krone to about 6 percent for the hedge position of the Swiss Franc. Hence, the increase in the systematic component of the liquidity costs is moderate compared to the case of widening or tightening market spreads.

Table 4.6: Jumps in market spreads (Case 4)
The calendar-spread static hedge is indicated by DEK and the strike-spread static hedge by CC. For each method, the results are given for the initial value of the static hedge for the up-and-out barrier call option, the average liquidity costs and the average liquidity costs as a percentage of the initial static hedge value. (*) denotes significant at 5% level and (**) denotes significant at 1% level.

Currency	Strike	DEK			CC		
		Value	Liq. costs	Liq. %	Value	Liq. costs	Liq. %
CHF	ITM	$3.99 \cdot 10^3$	210.8	5.3	$3.87 \cdot 10^3$	214.4	5.5
	ATM	$1.86 \cdot 10^3$	142.1	7.7	$1.77 \cdot 10^3$	142.3	8.1
	OTM	$6.09 \cdot 10^2$	68.6	11.3	$5.50 \cdot 10^2$	66.4	12.1
DKK	ITM	$8.25 \cdot 10^2$	24.4	3.0	$8.12 \cdot 10^2$	25.0	3.1
	ATM	$3.61 \cdot 10^2$	17.5	4.8	$3.51 \cdot 10^2$	17.7	5.0
	OTM	$1.09 \cdot 10^2$	8.8	8.1	$1.02 \cdot 10^2$	8.6	8.5
GBP	ITM	$7.15 \cdot 10^3$	165.7	2.3	$7.09 \cdot 10^3$	170.4	2.4
	ATM	$2.81 \cdot 10^3$	123.7	4.4	$2.76 \cdot 10^3$	125.5	4.5
	OTM	$7.42 \cdot 10^2$	60.0	8.1	$7.11 \cdot 10^2$	59.8	8.4
NOK	ITM	$6.81 \cdot 10^2$	61.0	9.0	$6.61 \cdot 10^2$	62.5	9.5
	ATM	$3.10 \cdot 10^2$	41.4	13.4	$2.95 \cdot 10^2$	41.9	14.2
	OTM	$9.95 \cdot 10^1$	20.5	20.6	$8.98 \cdot 10^1$	20.2	22.5

Table 4.7: Systematic liquidity costs
The calendar-spread static hedge is indicated by DEK and the strike-spread static hedge by CC. The systematic liquidity percentage of the total average liquidity costs are given below.

Currency	Market spreads	DEK			CC		
		ITM	ATM	OTM	ITM	ATM	OTM
CHF	Case 1	43.3	43.4	43.4	43.3	43.4	43.4
	Case 2	71.6	72.7	73.8	71.9	73.0	74.2
	Case 3	67.4	65.9	64.1	67.0	65.4	63.4
	Case 4	48.9	48.9	48.9	48.9	48.9	49.0
DKK	Case 1	82.7	82.7	82.7	82.7	82.7	82.7
	Case 2	94.1	94.4	94.6	94.2	94.4	94.7
	Case 3	92.7	92.2	91.7	92.6	92.1	91.5
	Case 4	85.6	85.6	85.6	85.6	85.6	85.6
GBP	Case 1	25.5	25.5	25.5	25.5	25.5	25.5
	Case 2	53.3	54.3	55.3	53.6	54.6	55.7
	Case 3	47.7	46.5	45.1	47.3	46.1	44.4
	Case 4	30.0	30.0	30.0	30.0	30.0	30.0
NOK	Case 1	25.6	25.7	25.7	25.7	25.7	25.7
	Case 2	54.1	55.4	56.9	54.5	55.9	57.4
	Case 3	47.2	45.2	43.0	46.6	44.6	42.0
	Case 4	30.1	30.1	30.1	30.1	30.1	30.1

4.6 Conclusions

The purpose of this paper is to examine the liquidity of the barrier option. Measuring the liquidity of (OTC) barrier options by the bid-ask spread can be a serious challenge, because (OTC) barrier option quotes are not widely available. The static hedging approach in this paper provides a solution to this problem, since the liquidity costs can be measured by only using the bid and ask quotes of the underlying asset prices in the framework of the Çetin, Jarrow, Protter and Warachka (2006) model.

Moreover, the approach in this paper allows for explicit computation of idiosyncratic and systematic liquidity costs, which can be relevant for financial institutions. As discussed in BIS (2006), financial institutions find it challenging to prepare for systematic liquidity risks.

Finally, the simulation study provides insight into the way time-varying market liquidity characteristics, such as widening, tightening market spreads and jumps in market spreads, can affect static barrier option hedging.

Barrier options are often traded on foreign exchange markets and four currencies with different liquidity characteristics have been examined in this paper. The main results are as follows. First, the results show that the liquidity of the Danish Krone, Swiss Franc, British Pound and the Norwegian Krone depend significantly on the market liquidity represented by the liquidity of the Euro. The most liquid currency is the Danish Krone, which has the smallest average liquidity costs and is highly dependent on the market liquidity. The Norwegian Krone is relatively the most illiquid currency and this illiquidity mainly arises from the idiosyncratic effect in the model. The liquidity costs of a static barrier option hedge for an illiquid currency can already be substantial in calm market conditions (benchmark case). For example, the liquidity costs of the static hedge for the Norwegian Krone vary from 8.9 percent for an ITM barrier option to 21.2 percent for an OTM barrier option in the benchmark.

Second, the simulation results reveal that the impact of time-varying market liquidity on static barrier option hedges can increase rapidly for different market scenarios. Widening market spreads, often associated with financial crises, cause the largest increase in liquidity costs compared to tightening market spreads and temporary jumps in market spreads. The results in this study show that the average liquidity cost percentage of a static hedge corresponding to very liquid currencies in tranquil market conditions can increase by about factor of 3 when market spreads become wider over time, if the liquidity of the currency is also highly dependent on the market liquidity, as in the case of the Danish Krone. As market spreads become wider over time, the systematic liquidity cost percentages approximately double for the British Pound and the Norwegian Krone, increasing from around 25 percent in the benchmark to more than 50 percent. Jumps in market bid-ask spreads can have a relatively high impact on liquidity costs of the static hedge for an OTM barrier option as well, considering the infrequent occurrence of jumps during the option life.

Third, liquidity risk of the static hedge also depends considerably on the moneyness level of the barrier option. Overall, the percentage impact of liquidity costs are relatively the highest for the static hedge of an initially OTM barrier option. This is similar to the result that Çetin, Jarrow, Protter and Warachka (2006) find for standard European options. The impact of widening market spreads has the highest systematic percentage liquidity impact on the OTM barrier option, whereas initially high market spreads have the largest systematic percentage impact on the static hedge for the ITM barrier option.

So, to determine the liquidity risk of a static barrier option hedge in the foreign exchange market, it is not sufficient to only consider the recent liquidity of the currency, but the dependency on the market liquidity and the moneyness level of the barrier option are very important as well. The liquidity costs of the static barrier option hedges can increase substantially due to time-varying market liquidity.

4.A Appendix

Barrier option formula

Let $S(t,0)$ be the foreign exchange mid rate at time t, K be the strike price, T be the option maturity, B the barrier level, σ the volatility, r the interest rate, r_f the foreign interest rate, and $\Phi(\cdot)$ the normal cumulative density function. The Black-Scholes call option formula $C(S(t,0),t|K,T)$ is given by

$$\begin{aligned}
C(S(t,0),t|K,T) &= S(t,0)e^{-r_f(T-t)}\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2), \quad \text{where} \\
d_1 &= \frac{\ln(S(t,0)/K) + (r - r_f + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}, \\
d_2 &= \frac{\ln(S(t,0)/K) + (r - r_f - \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}} = d_1 - \sigma\sqrt{T-t}.
\end{aligned}$$

The closed-form formula for the up-and-out barrier call $C_{uo}(S,t|K,B,T)$ is as follows

$$\begin{aligned}
C_{uo}(S(t,0),t|K,B,T) = \; & S(t,0) \quad e^{-r_f(T-t)}\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2) \\
& - \quad S(t,0)\Phi(x_1)e^{-r_f(T-t)} + Ke^{-r(T-t)}\Phi(x_1 - \sigma\sqrt{T-t}) \\
& + \quad S(t,0)e^{-r_f(T-t)}(B/S(t,0))^{2\lambda}[\Phi(-y) - \Phi(-y_1)] \\
& - \quad Ke^{-r(T-t)}(B/S(t,0))^{2\lambda-2}[\Phi(-y+\sigma\sqrt{T-t}) \\
& - \quad \Phi(-y_1+\sigma\sqrt{T-t})]], \qquad \text{where} \\
\lambda &= \frac{r - r_f + \sigma^2/2}{\sigma^2}, \\
y &= \frac{\ln(B^2/(S(t,0)K))}{\sigma\sqrt{T-t}} + \lambda\sigma\sqrt{T-t}, \\
x_1 &= \frac{\ln(S(t,0)/B)}{\sigma\sqrt{T-t}} + \lambda\sigma\sqrt{T-t}, \\
y_1 &= \frac{\ln(B/S(t,0))}{\sigma\sqrt{T-t}} + \lambda\sigma\sqrt{T-t}.
\end{aligned}$$

For more details see Hull (2006) and Zhang (1998).

Chapter 5

Covariance Shrinkage: Less is More?

5.1 Introduction

The framework of Basel II of the Bank for International Settlements requires banks to further improve risk management. A popular risk measure is value-at-risk (VaR), but it is often said that VaR cannot fully capture the effects of market risks such as sudden shifts in correlations or non-normality in the asset returns. In this paper we focus on different correlations specifications for asset returns and the resulting effects on VaR computations. To compute VaR for a portfolio with multiple assets one needs accurate estimates of the corresponding covariance matrices of these assets. Covariance matrices can be estimated by a great number of different multivariate GARCH models have been proposed to model covariance matrices, see Bauwens, Laurent and Rombouts (2003) for a survey. However, many multivariate GARCH models have a number of parameters that increases considerably when the number of assets increase or have other disadvantages. Therefore, parsimonious models such as the industry standard RiskMetrics^{TM} (see J.P. Morgan (1996)) are preferred for applications with many assets. Recently, Ledoit and Wolf (2003, 2004) have proposed a shrinkage method to estimate the covariance matrix of stock returns as an alternative to the sample covariance matrix. They combine the advantages of different covariance matrices using a weighting factor (the shrinkage constant).

Studies of Ledoit and Wolf (2003,2004) have shown that applying shrinkage techniques to obtain estimates of covariance matrix has many advantages such as computational simplicity and the estimated covariance matrices are always positive semi-definite. Moreover, Ledoit and Wolf (2003) show promising results using monthly stock data for the NYSE and AMEX stock returns, where the shrinkage model (with shrinkage target as single factor model) outperforms many alternatives as industry factor models and principal components approach in a mean-variance framework. Jagannathan and Ma (2003) have also found that this shrinkage model performs relatively well in large portfolios. Shrinkage techniques have also been applied by Jorion (1986) on the mean level using a different loss function specification and he shows that portfolio selection procedures based on sample means can be improved considerably.

These empirical studies have shown that applying shrinkage techniques to estimate mean or covariance matrices of returns can render considerable gains in portfolio optimization applications compared to alternative methods. In this paper the performance of the shrinkage estimation method for covariance matrices proposed in Ledoit and Wolf (2004) will be examined. They have applied shrinkage estimation techniques to US equity data for portfolio management using a constant-correlation model as shrinkage target and have found that the shrinkage method improves results in terms of higher (average) mean and lower standard deviations of excess returns compared to alternative models. However, recent studies such as Longin and Solnik (1995), Engle (2002), Ang and Chen (2002), Goetzmann, Li and Rouwenhorst (2005) and Connolly, Stivers and Sun (2005) have shown that correlations of asset returns are time varying. Thus, using a constant-correlation model to estimate the covariance matrices could result in serious misspecification. In empirical applications it is difficult to determine the accuracy of covariance matrix estimates, since the true covariance matrices are unknown. In order to see to what extent misspecification could affect the estimation results of covariance matrices and the resulting effects on VaR computations, this paper considers simulation experiments using several correlation specifications. In this paper the performance of the shrinkage estimator of Ledoit and Wolf (2004) will be compared to RiskMetrics^{TM} in a Monte Carlo simulation setting, where the true correlations are known and hence comparison between models can be performed using the known true correlations. In Section 5.2 the techniques used in this paper are discussed. Section

5.3 proceeds with the discussion of the simulation results obtained using zero correlations, constant non-zero correlations and time-varying correlations. Finally, Section 5.4 concludes.

5.2 Estimation methods

Consider the sample covariance matrix S of the stock returns, which has the advantages of computational simplicity and is an unbiased estimate of the true covariance matrix. A disadvantage of the sample covariance matrix is that it involves a lot of estimation error, especially when the number of assets is relatively large compared to the number of return observations available for each asset. The alternative approach would be to use a highly structured covariance matrix F, that has less estimation error but could have a bias. Ledoit and Wolf (2003) proposed a way to make an optimal trade-off between estimation error and bias by combining S and F using a weighting factor $\delta \in [0, 1]$, such that $\delta F + (1 - \delta)S$ becomes the new covariance matrix. This technique is called shrinkage and has been introduced in statistics by Stein (1955). Ledoit and Wolf (2004) have chosen for the constant correlation model (see Elton and Gruber (1973)) as the structured covariance matrix F. The approach taken by Ledoit and Wolf is as follows.

The crucial step in the shrinkage method is how to determine the optimal shrinkage constant δ. The optimal δ should minimize the distance between the shrinkage covariance matrix and the true covariance matrix measured by the Frobenius norm (defined by $\| Z \|^2 = \sum_{i=1}^{N} \sum_{j=1}^{N} z_{ij}^2$) . Consider the following loss function

$$\mathrm{L}(\delta) = \| \delta F + (1 - \delta)S - \Sigma \|^2 . \tag{5.1}$$

The optimal δ should minimize the expected loss function, i.e. $\mathbb{E}(L(\delta))$. Let N be the number of assets, let $\{s_{ij}\}_{i,j=1}^{N}$ be the elements of the sample covariance matrix S. For the constant correlation model, $\{\phi_{ij}\}_{i,j=1}^{N}$ are elements of the population covariance matrix Φ and $\{f_{ij}\}_{i,j=1}^{N}$ are the elements corresponding to the structured covariance matrix F. The diagonal elements of F are equal to s_{ii}, whereas the off-diagonal elements are given by $\bar{r}\sqrt{s_{ii}s_{jj}}$ ($\bar{r}$ is the average of the sample correlations). Ledoit and Wolf

(2003) show that the optimal δ is given by κ/T, such that

$$\begin{aligned}
\kappa &= \frac{\pi - \rho}{\gamma}, \\
\text{where} \qquad \pi &= \sum_{i=1}^{N}\sum_{j=1}^{N} \text{AVar}\left[\sqrt{T} s_{ij}\right], \\
\rho &= \sum_{i=1}^{N}\sum_{j=1}^{N} \text{ACov}\left[\sqrt{T} f_{ij}, \sqrt{T} s_{ij}\right], \\
\gamma &= \sum_{i=1}^{N}\sum_{j=1}^{N} (\phi_{ij} - \sigma_{ij})^2.
\end{aligned}$$

The estimated shrinkage constant can be obtained by

$$\hat{\delta}_* = \max\left\{0, \min\{\frac{\hat{\kappa}}{T}, 1\}\right\}. \tag{5.2}$$

For the consistent estimate $\hat{\kappa}$ and more details see Ledoit and Wolf (2004). The results of the shrinkage method will be compared to the RiskMetrics^{TM} model. Let r_t be the $1 \times N$ vector of asset returns at time t ($t = 1, \ldots, T$) and let H_t be the conditional covariance matrix at time t. The RiskMetrics^{TM} model is given as

$$H_t = \lambda H_{t-1} + (1 - \lambda) r'_{t-1} r_{t-1},$$

where λ is taken to be 0.94 for daily observations.

5.3 Monte Carlo Simulation

Since the introduction of the ARCH models by Engle (1982) and GARCH models by Bollerslev (1986), many studies have shown that the low order GARCH model captures the empirical features of asset returns very well. Hansen and Lunde (2001) have examined 330 different ARCH-type models and have found that for exchange rates the standard GARCH(1,1) model cannot be outperformed by other models. However, for equity returns the models that account for asymmetric effects clearly performed better. Therefore, we will use the following asymmetric GARCH specification of Glosten, Jagannathan and Runkle (1993) and Zakoian (1994) for our simulations. This model can be interpreted as an unrestricted specification of the GARCH(1,1) model. Let the asset returns $r_t = \mu + \epsilon_t$ $(t = 1, \ldots, T)$ be normally distributed with mean μ and the conditional variance given by

$$h_t = \omega + \beta h_{t-1} + \alpha_+ \epsilon^2_{+,t-1} + \alpha_- \epsilon^2_{-,t-1}, \tag{5.3}$$

where $\epsilon_{-,t}$ contains only the negative values of ϵ_t and the positive values are replaced with zeroes. Analogously, $\epsilon_{+,t}$ only contains the positive values of ϵ_t. Henceforth, the returns are assumed to be mean zero for simplicity.

The returns of three assets will be simulated using multivariate normally distributed errors with several different correlation specifications. On the volatility level, both the effect of persistence as well as asymmetry will be examined by specifying two different sets of volatility data generating processes for $t = 1, \ldots, T$ and $i, j = 1, \ldots, N$:

DGP 1:

$$\begin{aligned} h_{1,t} &= 0.01 + 0.20 h_{1,t-1} + 0.15 r^2_{1+,t-1} + 0.15 r^2_{1-,t-1}, \\ h_{2,t} &= 0.01 + 0.40 h_{2,t-1} + 0.20 r^2_{2+,t-1} + 0.20 r^2_{2-,t-1}, \\ h_{3,t} &= 0.01 + 0.95 h_{3,t-1} + 0.02 r^2_{3+,t-1} + 0.02 r^2_{3-,t-1}, \end{aligned} \tag{5.4}$$

DGP 2:

$$\begin{aligned} h_{1,t} &= 0.01 + 0.20 h_{1,t-1} + 0.05 r^2_{1+,t-1} + 0.25 r^2_{1-,t-1}, \\ h_{2,t} &= 0.01 + 0.40 h_{2,t-1} + 0.30 r^2_{2+,t-1} + 0.10 r^2_{2-,t-1}, \\ h_{3,t} &= 0.01 + 0.95 h_{3,t-1} + 0.01 r^2_{3+,t-1} + 0.03 r^2_{3-,t-1}, \\ r_{i,t} &= \sqrt{h_{i,t}} \epsilon_{i,t}, \quad i = 1, 2, 3, \quad \text{and} \quad \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \\ \epsilon_{3,t} \end{pmatrix} \sim \mathcal{N}(0, R_t). \end{aligned} \tag{5.5}$$

DGP 1 basically amounts to standard GARCH(1,1) specifications, where different persistence levels are included to examine how this affects the results of the covariances. DGP 2 gives volatility series that exhibit asymmetric effects, since these are often observed in empirical studies for stock return data. As mentioned earlier, Ledoit and Wolf (2004) apply a constant correlation model to US stock returns. However, recent studies such as Longin and Solnik (1995), Engle (2002), Ang and Chen (2002), Goetzmann, Li and Rouwenhorst (2005) and Connolly, Stivers and Sun (2005) have shown that correlations of asset returns are time varying. Therefore, this paper considers several correlation specifications, where the simulation setting is partly similar as in Engle (2002), although he compares the RiskMetrics^{TM} model to other models than the shrinkage estimation model in this paper.

Figure 5.1: Correlations generated by RiskMetrics^{TM} model

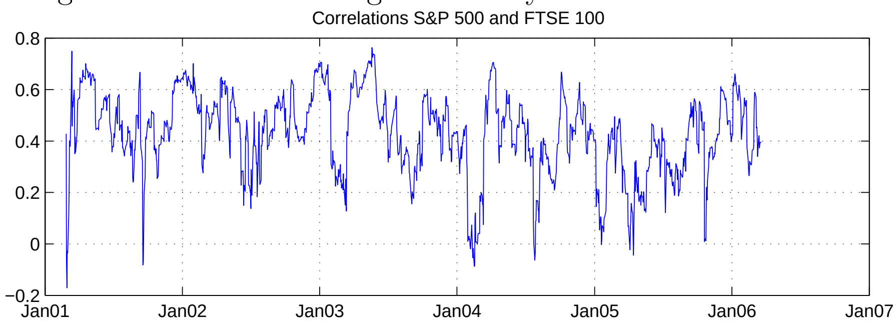

The following three correlation specifications will be considered: zero, constant (nonzero) and time-varying correlations. For a zero correlation experiment the matrix R is simply the identity matrix. Moreover, the following specifications of R_t will be considered for $t = 1, \ldots, T$ and $i, j = 1, \ldots, N$:

Constant correlations:

$$R = \begin{pmatrix} 1 & 0.8 & 0.1 \\ 0.8 & 1 & -0.3 \\ 0.1 & -0.3 & 1 \end{pmatrix}, \tag{5.6}$$

Time-varying correlations:

$$R_t(i,j) = \begin{cases} (0.5 + 0.4\cos(2\pi t/250)) \cdot R_{ij} & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases}, \tag{5.7}$$

where R_{ij} are the corresponding elements of R matrix (5.6). The specification of the time-varying correlations in (5.7) has been chosen as a simplified pattern that resembles the time-varying correlations computed using empirical data as, for example, in Figure 5.1. All the simulation experiments will be performed for the symmetric as well as the asymmetric volatility specifications (respectively, DGP 1 and DGP 2). First, the returns are simulated using zero correlation, then using constant correlations and finally using time-varying correlations, which amounts to six different simulation experiments. In each of the simulation experiments, the RiskMetrics^{TM} and shrinkage method using the constant correlation model as shrinkage target according to Ledoit and Wolf (2004)[1] are used to estimate covariance matrices in a rolling window framework for several window sizes. The comparison between the models under consideration will be twofold, using statistical as well as economic criteria. The accuracy of the estimation results will be assessed by statistical measures as the mean absolute error (henceforth, MAE) and the root mean squared error (RMSE) as well as by an economic measure as value-at-risk (VaR). In a simulation experiment, the true correlations are known in advance, hence the estimated correlations can be easily compared to the true correlations using the statistical loss functions MAE and RMSE. However, covariance (or correlation) models are often used in empirical applications, where the true values are unknown, but where economic measures such as VaR require accurate estimates of the covariance matrix. To compare the correlation models using VaR an equally weighted portfolio of the three simulated asset return series will be used and the VaR is computed as

$$\mathrm{VaR}_t = z_\alpha \sqrt{w'\Sigma_t w}, \quad t = 1, \ldots, T. \tag{5.8}$$

where α is the confidence level corresponding to the VaR (z is the standard normal cdf value) and the initial wealth has been normalized to one.

For each simulation, the VaR for each time period will be computed using the estimated covariance matrix. Hence, each simulation gives a series of T VaR observations and the number of VaR exceedances in the sample will be counted. Provided that the variance and covariance estimates are accurate, the number of violations of the 95% VaR level should be very close to 5% of sample size T, since the standardized returns ϵ_t are normally distributed by construction. In the next sections, both volatility data generating processes will be used in the different correlation experiments to see whether asymmetry effects in volatility will affect the results for the correlations. Moreover, the

[1] I have used the program of the shrinkage covariance matrix estimator provided at the author's website 'www.ledoit.net'. For more information, see Ledoit and Wolf (2004).

different correlation experiments will be used to assess the performance of the models under consideration.

5.3.1 Zero correlations

The first Monte Carlo simulation experiment will be performed using multivariate normally distributed errors generated with zero correlations across assets. If the individual volatilities are not well captured by the covariance model, it is very likely that the remaining volatility effects could be incorrectly interpreted as correlation effects. The results of the simulation using zero correlations between assets can give more insight in the effects that different volatility specifications have on the correlations estimates, even though the true correlations are zero. Hence, using the results in this section as a benchmark for the subsequent experiments where correlations are nonzero, can help to separate the volatility effects from the correlation effects on the estimation results. The results of this experiment for both volatility DGP 1 and 2 are given in Table 5.1 to 5.2. The outline of the discussion of the results will be as follows. The results obtained from DGP 1 will be discussed concerning different volatility persistence across assets, different sample window sizes for RiskMetrics^{TM} and the shrinkage model[2]. Judging by the large MAE and RMSE values in Table 5.1 the volatility effects have a considerable impact on the correlation estimates for the RiskMetrics^{TM} model. Hence, the volatility effects can cause correlation estimates to be nonzero, even though the true correlations are zero. For each experiment the sample size T is 1000 and the simulations are repeated 1000 times. Table 5.1 shows that for the RiskMetrics^{TM} model the average of the MAE values of the replications is relatively stable across different window sizes 25, 50,100 and 250 for each asset. For the shrinkage method increasing window size does increase the accuracy of the correlation estimates considerably, resulting in declining MAE for larger window sizes. So, for zero true correlations, the shrinkage estimation method outperforms the RiskMetrics^{TM} method in statistical terms, since it has smaller MAE and RMSE values for the larger window sizes considered.

---

[2]For RiskMetrics^{TM}, the window size only affects the initial covariance matrix estimation at the beginning of the sample for each simulation. Subsequent covariance matrices are obtained by updating using the return shocks and the previous estimate of the covariance matrix. For the shrinkage method, the optimal shrinkage factor is estimated using a rolling sample for different window sizes.

Regarding the effects of different volatility persistence levels, Table 5.1 shows that the errors for ρ_{21} are smaller than for ρ_{31}, and the errors of the latter are in turn smaller than those of ρ_{32}. This result is consistent for all window sizes and holds both for the RiskMetrics^{TM} and shrinkage method. Since all correlations have true value of zero, this effect should be ascribed to the volatility effects. The volatility specifications $h_{1,t}$, $h_{2,t}$ and $h_{3,t}$ have increasing persistence, indicating that mean reversion to the long run volatility level is decreasing. Apparently, the correlation between the two assets with the least persistent volatility series $h_{1,t}$ and $h_{2,t}$ have relatively smaller errors than the correlation between those of the two more persistent volatility series $h_{2,t}$ and $h_{3,t}$, though the differences are not extremely large. Table 5.2 shows that the shrinkage method gives VaR results that are more in line with what should be expected according to theory using window sizes larger than 50, whereas the RiskMetrics^{TM} model gives relatively more VaR violations.

Table 5.1: MAE and RMSE for zero correlations
This table gives the average MAE and RMSE over 1000 simulations for each correlation series, where the data is generated using DGP 1 and 2 with zero correlations between the multivariate normal errors of the asset returns. The window sizes used for estimation of the covariance matrix are respectively 25, 50, 100 and 250 observations.

	RiskMetrics^{TM}			Shrinkage		
	ρ_{21}	ρ_{32}	ρ_{31}	ρ_{21}	ρ_{32}	ρ_{31}
True ρ	0	0	0	0	0	0
Symmetric volatility (DGP 1)						
MAE25	0.134	0.140	0.139	0.140	0.143	0.142
MAE50	0.134	0.139	0.138	0.093	0.097	0.095
MAE100	0.133	0.138	0.138	0.063	0.066	0.065
MAE250	0.132	0.137	0.137	0.039	0.041	0.039
RMSE25	0.171	0.175	0.174	0.178	0.181	0.179
RMSE50	0.170	0.174	0.173	0.119	0.122	0.121
RMSE100	0.168	0.173	0.172	0.080	0.083	0.082
RMSE250	0.167	0.171	0.170	0.048	0.050	0.049
Asymmetric volatility (DGP 2)						
MAE25	0.133	0.139	0.138	0.140	0.145	0.143
MAE50	0.132	0.139	0.138	0.094	0.098	0.096
MAE100	0.132	0.138	0.137	0.064	0.067	0.065
MAE250	0.131	0.137	0.136	0.039	0.041	0.040
RMSE25	0.170	0.175	0.174	0.178	0.182	0.180
RMSE50	0.168	0.174	0.172	0.119	0.123	0.121
RMSE100	0.167	0.172	0.171	0.081	0.084	0.082
RMSE250	0.166	0.171	0.169	0.048	0.050	0.049

Table 5.2: VaR simulation results with zero correlations
This table gives the theoretical violation rates of 0.10, 0.05, 0.01 of respectively the 90%, 95% and 99% VaR level, and the average VaR exceedance rates (and corresponding standard deviations) across 1000 simulations. The sample size of each simulation is $T = 1000$.

	DGP 1				DGP 2			
	RiskMetrics^{TM}		Shrinkage		RiskMetrics^{TM}		Shrinkage	
	Rate	S.d. ($\cdot 10^{-4}$)	Rate	S.d. ($\cdot 10^{-4}$)	Rate	S.d. ($\cdot 10^{-4}$)	Rate	S.d. ($\cdot 10^{-4}$)
90% VaR	0.100	–	0.100	–	0.100	–	0.100	–
VaR25	0.105	2.64	0.109	2.71	0.105	2.67	0.108	2.75
VaR50	0.105	2.64	0.103	2.63	0.104	2.67	0.102	2.71
VaR100	0.105	2.63	0.100	2.62	0.104	2.66	0.099	2.68
VaR250	0.104	2.63	0.098	2.68	0.103	2.65	0.097	2.78
95% VaR	0.050	–	0.050	–	0.050	–	0.050	–
VaR25	0.057	1.95	0.060	2.02	0.057	1.98	0.060	2.03
VaR50	0.057	1.95	0.055	1.94	0.057	1.97	0.054	1.95
VaR100	0.057	1.95	0.052	1.95	0.056	1.96	0.052	1.91
VaR250	0.056	1.93	0.050	1.93	0.056	1.94	0.050	1.95
99% VaR	0.010	–	0.010	–	0.010	–	0.010	–
VaR25	0.016	1.10	0.018	1.18	0.017	1.12	0.018	1.19
VaR50	0.016	1.10	0.014	1.07	0.017	1.11	0.015	1.10
VaR100	0.016	1.09	0.013	1.03	0.016	1.10	0.013	1.05
VaR250	0.015	1.06	0.012	1.01	0.016	1.08	0.013	1.02

5.3.2 Constant correlations

The simulations in this section address the case of asset returns with nonzero, constant correlations. The results are obtained using data generated by DGP 1 and DGP 2 for volatility and (5.6) for the corresponding correlations. Table 5.3 shows that increasing the window size decreases the errors measured by MAE and RMSE for the shrinkage method, but the errors for RiskMetrics^{TM} remain relatively unaffected by differences in window size. Also, more persistence and asymmetric effects lead to larger errors. These results are consistent with the results obtained for zero correlations in the previous section. A striking difference in this section is that the MAE of the ρ_{21} series for RiskMetrics^{TM} is much smaller than for the other correlation series, see the second column in Table 5.3. Thus, RiskMetrics^{TM} appears to give more accurate correlation estimates, when the true correlation value is high and the volatility is symmetric. This could be explained by the fact that RiskMetrics^{TM} models the volatilities using a common parameter and hence it is implicitly assumed that these processes are the same across all individual assets. When the returns between two assets are highly correlated, this assumption is more likely to be true. Therefore, RiskMetrics^{TM} appears to give more accurate correlation estimates for high true correlations, though this improvement is diminished when the volatilities exhibit asymmetric effects (see Table 5.3). For the shrinkage method, the correlation estimates deteriorate for nonzero true correlations as indicated by the higher MAE and RMSE values for all window sizes. Compared to the results of the previous section, the errors have increased for asymmetric volatility in Table 5.3, in particular for the estimates of ρ_{21} given in the fifth column of Table 5.3. Unlike for the other correlation series, this error cannot be reduced easily by increasing the window size. The average RiskMetrics^{TM} VaR exceedance rates are slightly lower for the significance levels 90% and 95% compared to those of the zero correlations, but the average percentile for the 99% VaR level is a bit higher. This result has also been found for the shrinkage method. However, the shrinkage method has a smaller average of exceedances rates in most cases compared to RiskMetrics^{TM}.

5.3.3 Time-varying correlations

Using time-varying correlations the results change considerably compared to earlier results. Since the shrinkage method is developed under the assumption of constant correlations, there is misspecification of the model in this experiment. Many recent studies mentioned earlier have shown that correlations are not constant. Yet, in several studies the constant correlation model has been empirically applied as mentioned earlier. Hence, analyzing simulated asset return data using time-varying correlations can shed some light on the severeness of the model misspecification effects for the accuracy of the covariance estimates.

For constant correlations the shrinkage method outperformed the RiskMetrics^{TM} model, but in this section RiskMetrics^{TM} is clearly preferred to the shrinkage method according to MAE and RMSE in Table 5.5. Both statistical criteria are relatively large compared to the case of zero and constant correlations. For the shrinkage method, the window size is optimal for around 50 to 100 observations and the other window sizes have larger errors. Choosing a too small window size could result in the fact that the

Table 5.3: MAE and RMSE for DGP 1 and constant nonzero correlations

This table gives the average MAE and RMSE over 1000 simulations for each correlation series, where the data is generated using DGP 1 and 2 with constant nonzero correlations between the multivariate normal errors of the asset returns. The window sizes used for estimation of the covariance matrix are respectively 25, 50, 100 and 250 observations.

	RiskMetrics^{TM}			Shrinkage		
	ρ_{21}	ρ_{32}	ρ_{31}	ρ_{21}	ρ_{32}	ρ_{31}
True ρ	0.8	-0.3	0.1	0.8	-0.3	0.1
Symmetric volatility (DGP 1)						
MAE25	0.083	0.128	0.138	0.115	0.159	0.157
MAE50	0.083	0.128	0.137	0.093	0.112	0.111
MAE100	0.083	0.127	0.137	0.078	0.079	0.078
MAE250	0.082	0.126	0.136	0.066	0.052	0.050
RMSE25	0.109	0.162	0.173	0.148	0.200	0.195
RMSE50	0.108	0.161	0.172	0.114	0.140	0.137
RMSE100	0.107	0.160	0.170	0.091	0.099	0.097
RMSE250	0.107	0.158	0.169	0.072	0.063	0.061
Asymmetric volatility (DGP 2)						
MAE25	0.128	0.129	0.137	0.161	0.158	0.157
MAE50	0.128	0.128	0.137	0.153	0.112	0.111
MAE100	0.128	0.128	0.136	0.147	0.081	0.079
MAE250	0.128	0.127	0.135	0.141	0.054	0.051
RMSE25	0.153	0.163	0.173	0.192	0.199	0.195
RMSE50	0.153	0.162	0.171	0.171	0.141	0.137
RMSE100	0.153	0.161	0.170	0.157	0.100	0.097
RMSE250	0.152	0.159	0.168	0.145	0.065	0.062

shrinkage method is subject to sample estimation error, whereas a too large window cannot capture the time variation in the correlation. If asymmetric volatility effects are considered, the errors for the shrinkage method become very large.

For the VaR results, RiskMetrics^{TM} has comparable performance for different volatility and correlation specifications. For constant correlations, the shrinkage method outperforms RiskMetrics^{TM}, but the performance for time-varying correlations deteriorates considerably as measured by statistical criteria, although this does not seem to affect the results for the VaR violations.

5.3.4 Shrinkage factor analysis

The shrinkage factor δ is very important for the model of Ledoit and Wolf (2004), since it determines the degree of structure included in the estimation of the covariance

Table 5.4: VaR simulation results with constant nonzero correlations
This table gives the theoretical violation rates of 0.10, 0.05, 0.01 of respectively the 90%, 95% and 99% VaR level, and the average VaR exceedance rates (and corresponding standard deviations) across 1000 simulations. The sample size of each simulation is $T = 1000$.

	DGP 1				DGP 2			
	RiskMetrics^{TM}		Shrinkage		RiskMetrics^{TM}		Shrinkage	
	Rate	S.d. $(\cdot 10^{-4})$	Rate	S.d. $(\cdot 10^{-4})$	Rate	S.d. $(\cdot 10^{-4})$	Rate	S.d. $(\cdot 10^{-4})$
90% VaR	0.100	–	0.100	–	0.100	–	0.100	–
VaR25	0.101	2.56	0.105	2.66	0.101	2.58	0.105	2.63
VaR50	0.101	2.57	0.098	2.64	0.101	2.58	0.098	2.58
VaR100	0.101	2.57	0.094	2.64	0.101	2.58	0.094	2.64
VaR250	0.100	2.56	0.092	2.71	0.100	2.58	0.092	2.71
95% VaR	0.050	–	0.050	–	0.050	–	0.050	–
VaR25	0.056	1.91	0.060	2.00	0.056	1.97	0.060	2.05
VaR50	0.056	1.91	0.054	1.92	0.056	1.96	0.054	1.95
VaR100	0.056	1.90	0.051	1.91	0.056	1.97	0.050	1.92
VaR250	0.055	1.88	0.048	1.91	0.055	1.94	0.049	1.98
99% VaR	0.010	–	0.010	–	0.010	–	0.010	–
VaR25	0.018	1.19	0.020	1.27	0.018	1.18	0.020	1.32
VaR50	0.018	1.18	0.016	1.16	0.018	1.17	0.016	1.18
VaR100	0.018	1.17	0.015	1.09	0.018	1.15	0.015	1.09
VaR250	0.017	1.15	0.014	1.08	0.017	1.13	0.014	1.08

matrix. If δ is close to zero, the resulting covariance matrix will resemble the sample covariance matrix very much and there is not much structure involved for the estimation of the covariance matrix. On the other hand, if the shrinkage parameter is close to unity, then the resulting covariance matrix is composed entirely by the structured covariance matrix F.

To examine the shrinkage factor more closely, Figures 5.2 to 5.8 given show the 25%-quantile, average and 75%-quantile of the shrinkage factor δ across the simulations for each time period of the sample. From Figures 5.2 and 5.3 we can see that in case of zero correlations, the optimal shrinkage factor estimates are very dispersed across the simulations. The average δ over the simulations is around 0.5, whereas the 25%-quantile is below or around 0.2 (for respectively a window size of 50 and 100) and the 75%-quantile is almost equal to 1 for most of the time periods. The graphs for symmetric volatility do not seem to be very different from the graphs of asymmetric volatility. However, as already concluded from the statistical results the choice of window size has a larger impact on the shrinkage method. The average shrinkage factor is slightly higher for the larger window size of 100 observations and the 25%- and

Table 5.5: MAE and RMSE for DGP 1 and time-varying correlations

This table gives the average MAE and RMSE over 1000 simulations for each correlation series, where the data is generated using DGP 1 and 2 with time-varying correlations between the multivariate normal errors of the asset returns. The window sizes used for estimation of the covariance matrix are respectively 25, 50, 100 and 250 observations.

	RiskMetrics^{TM}			Shrinkage		
	ρ_{21}	ρ_{32}	ρ_{31}	ρ_{21}	ρ_{32}	ρ_{31}
True ρ	–	–	–	–	–	–
Symmetric volatility (DGP 1)						
MAE25	0.142	0.137	0.139	0.160	0.160	0.150
MAE50	0.142	0.137	0.138	0.151	0.119	0.105
MAE100	0.142	0.136	0.138	0.205	0.106	0.077
MAE250	0.141	0.136	0.137	0.256	0.105	0.057
RMSE25	0.179	0.173	0.174	0.201	0.199	0.187
RMSE50	0.178	0.172	0.173	0.187	0.147	0.131
RMSE100	0.178	0.171	0.172	0.241	0.129	0.096
RMSE250	0.177	0.169	0.170	0.291	0.123	0.069
Asymmetric volatility (DGP 2)						
MAE25	0.149	0.137	0.138	0.169	0.160	0.150
MAE50	0.149	0.137	0.138	0.162	0.119	0.105
MAE100	0.149	0.136	0.137	0.208	0.106	0.077
MAE250	0.149	0.136	0.136	0.258	0.105	0.057
RMSE25	0.186	0.173	0.174	0.211	0.199	0.188
RMSE50	0.186	0.172	0.172	0.199	0.148	0.131
RMSE100	0.186	0.171	0.171	0.248	0.129	0.096
RMSE250	0.185	0.169	0.169	0.295	0.124	0.069

75%-quantiles indicate a larger interval than the quantiles for a window size of 100. The parameter γ in the numerator of κ is relatively small. Thus, for zero correlations the optimal shrinkage factor is dispersed across the interval from zero to one for the simulations considered.

For constant nonzero correlations, the values are very different for the optimal δ as shown in Figures 5.4 and 5.5. A notable difference of these pictures is that the optimal shrinkage parameter is very close to zero and the quantiles are much closer to the average δ than in case of zero correlations. Again, the asymmetric effects in volatility do not seem to affect the parameter estimates differently from symmetric volatility. The average of the shrinkage factor estimates is somewhat higher for window of 50 observations and the 25%- and 75%-quantiles indicate a larger interval than those for window size of 100 observations. Thus, for zero and constant nonzero correlations the larger window size gives estimates of shrinkage parameters that are more precise in

Table 5.6: VaR simulation results with time-varying correlations
This table gives the theoretical violation rates of 0.10, 0.05, 0.01 of respectively the 90%, 95% and 99% VaR level, and the average VaR exceedance rates (and corresponding standard deviations) across 1000 simulations. The sample size of each simulation is $T = 1000$.

	DGP 1				DGP 2			
	RiskMetrics^{TM}		Shrinkage		RiskMetrics^{TM}		Shrinkage	
	Rate	S.d. ($\cdot 10^{-4}$)	Rate	S.d. ($\cdot 10^{-4}$)	Rate	S.d. ($\cdot 10^{-4}$)	Rate	S.d. ($\cdot 10^{-4}$)
90% VaR	0.100	–	0.100	–	0.100	–	0.100	–
VaR25	0.104	2.67	0.108	2.70	0.103	2.57	0.107	2.63
VaR50	0.104	2.67	0.101	2.63	0.103	2.58	0.100	2.64
VaR100	0.104	2.66	0.098	2.60	0.103	2.57	0.097	2.66
VaR250	0.103	2.65	0.096	2.67	0.102	2.55	0.095	2.73
95% VaR	0.050	–	0.050	–	0.050	–	0.050	–
VaR25	0.057	1.89	0.060	1.97	0.057	1.93	0.060	2.06
VaR50	0.057	1.89	0.054	1.92	0.057	1.93	0.054	1.87
VaR100	0.056	1.89	0.051	1.93	0.056	1.93	0.051	1.88
VaR250	0.056	1.89	0.050	1.93	0.056	1.91	0.050	1.91
99% VaR	0.010	–	0.010	–	0.010	–	0.010	–
VaR25	0.017	1.17	0.018	1.21	0.017	1.17	0.018	1.23
VaR50	0.017	1.16	0.015	1.13	0.017	1.17	0.015	1.13
VaR100	0.017	1.15	0.014	1.09	0.017	1.15	0.014	1.08
VaR250	0.016	1.14	0.013	1.06	0.016	1.14	0.013	1.05

the sense that the 25%- and 75%-quantiles are closer to each other. When correlations are close to zero, the optimal shrinkage factor varies relatively greater, since the interval indicated by the 25%- to 75%-quantiles is quite large. The parameter γ in the numerator of κ is relatively small. When the correlations are constant and nonzero as specified in (5.6), the optimal shrinkage factor is very close to zero. Then, the resulting shrinkage covariance matrix estimation more or less amounts to simply taking the sample covariance matrix. Both results above can be found combined in the case of time-varying correlations of (5.7) (and are dependent on the window size), when comparing Figure 5.6 to Figure 5.7 and 5.8. For the time period with high nonzero correlations the shrinkage factor δ estimates are very close to zero and at the same time the dispersion of the optimal δ value decreases. The time periods for which the correlations are nearly zero coincide with the periods for which the shrinkage factor is large on average, but at the same time has a large range for 25%- and 75%-quantile. This results in an almost countercyclical pattern between the correlations and the optimal shrinkage factor estimates.

Figure 5.2: The 25%-quantile, average and 75%-quantile of the shrinkage factor δ simulations over time for DGP 1 and zero correlations

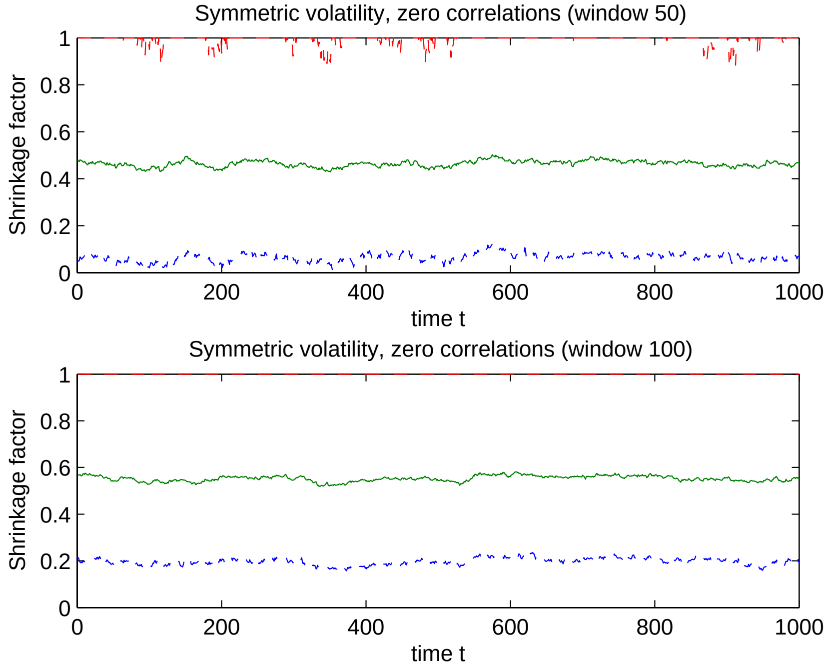

Figure 5.3: The 25%-quantile, average and 75%-quantile of the shrinkage factor δ simulations over time for DGP 2 and zero correlations

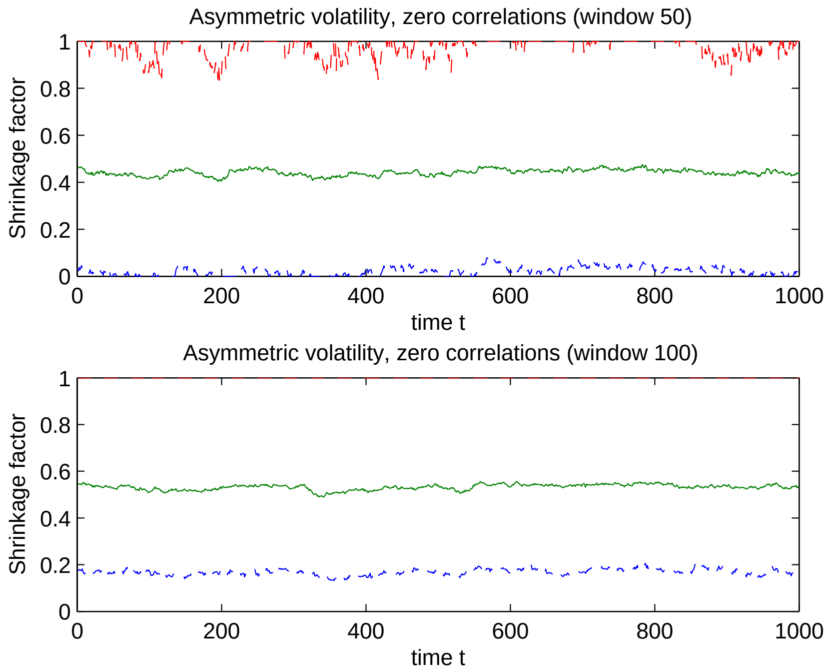

Figure 5.4: The 25%-quantile, average and 75%-quantile of the shrinkage factor δ simulations over time for DGP 1 and constant nonzero correlations

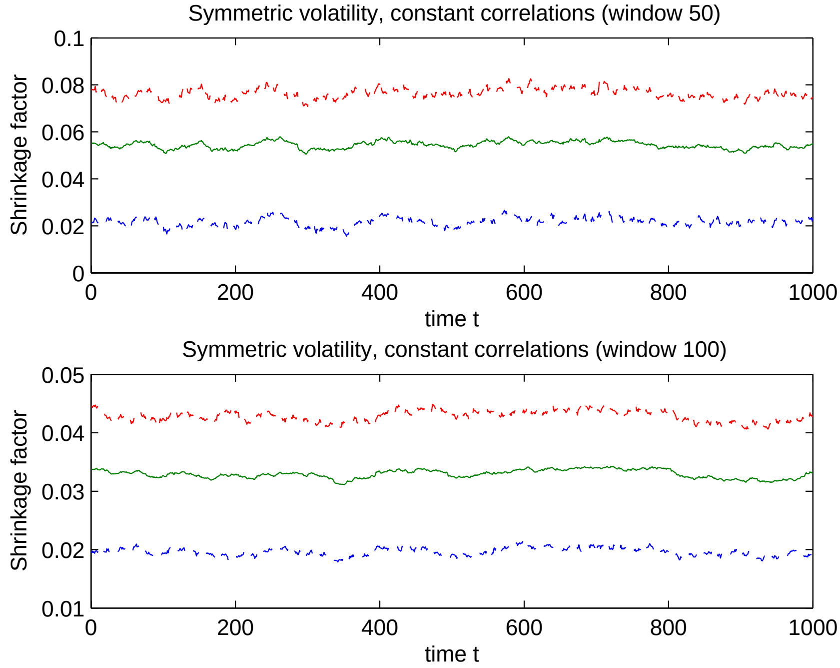

Figure 5.5: The 25%-quantile, average and 75%-quantile of the shrinkage factor δ simulations over time for DGP 2 and constant nonzero correlations

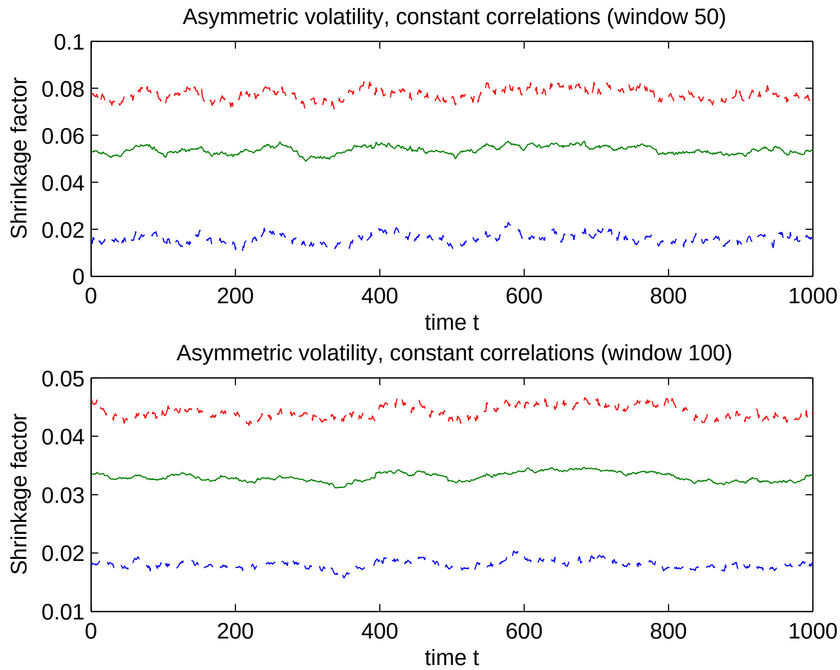

Figure 5.6: True time-varying correlations specified by (5.7)

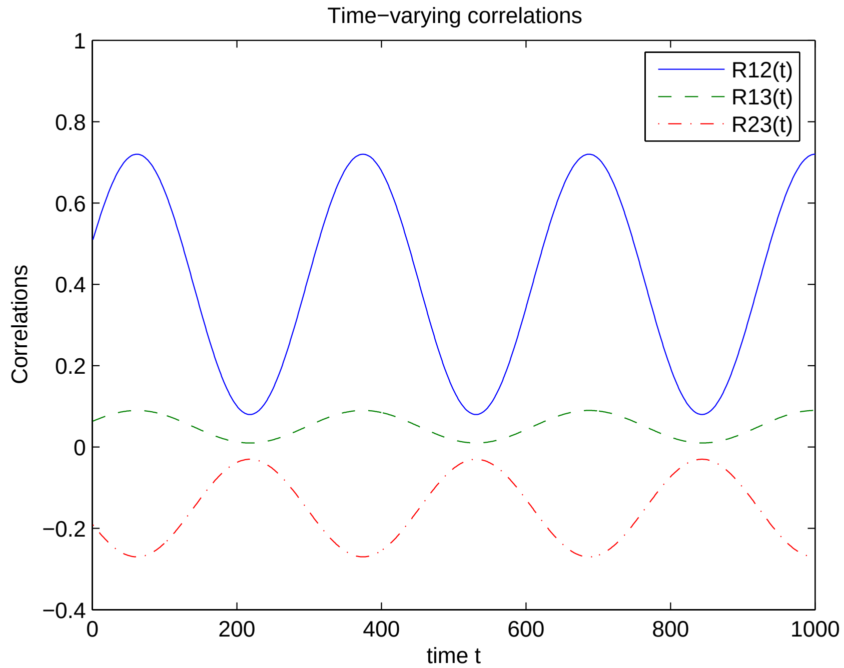

Figure 5.7: The 25%-quantile, average and 75%-quantile of the shrinkage factor δ simulations over time for DGP 1 and time-varying correlations

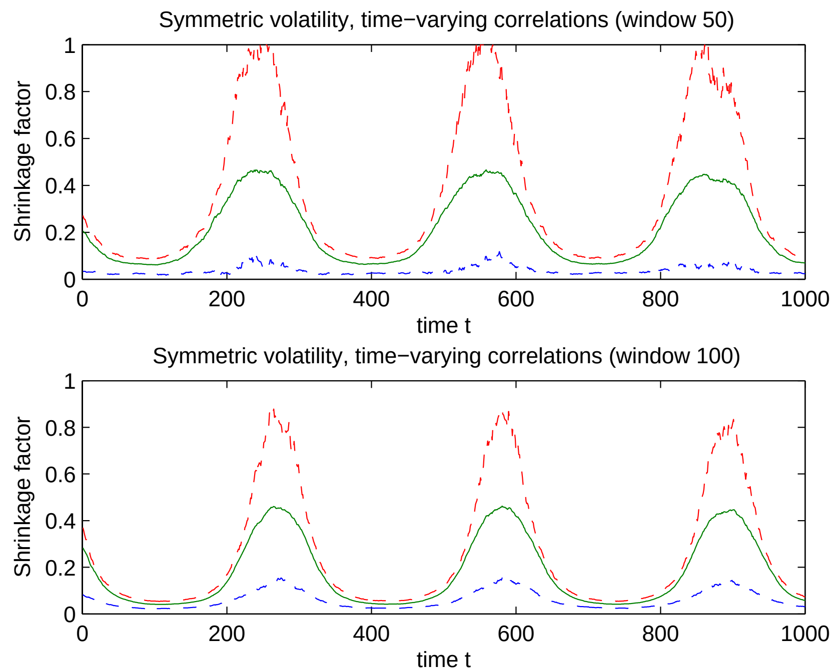

Figure 5.8: The 25%-quantile, average and 75%-quantile of the shrinkage factor δ simulations over time for DGP 2 and time-varying correlations

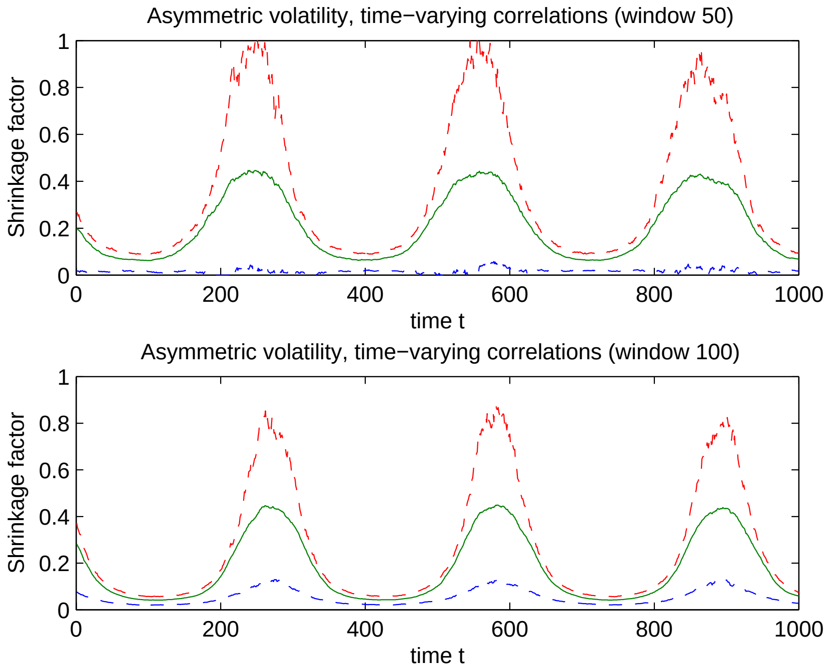

5.4 Conclusions

In this paper a Monte Carlo simulation experiment has been performed to compare the shrinkage method of Ledoit and Wolf (2004) to the industry standard RiskMetrics^{TM} to estimate covariance matrices. Different simulation settings have been employed, where symmetric and asymmetric volatility specifications combined with zero, constant nonzero and time-varying correlations between the asset returns have been included in this study. The following results have been found. First, applying shrinkage outperforms the RiskMetrics^{TM} for window sizes of 50 observations and larger in terms of lower MAE and RMSE and having VaR violation rates that are more in line with the theoretical rates for constant (zero and nonzero) correlations. These results are robust for different volatility persistence levels and for asymmetric volatility effects. It appears that the choice of window size has a much larger effect on the estimation results for the shrinkage method than volatility effects. The statistical loss functions do not show much difference in accuracy of correlation estimates between symmetric and asymmetric volatility specifications in case of zero correlations. However, for constant nonzero correlations the asymmetric volatility experiment renders higher MAE and RMSE for the correlation estimates corresponding to asset returns with high positive true correlations.

For time-varying correlations the performance of the shrinkage method changes drastically and RiskMetrics^{TM} is much better according to the statistical criteria. However, this does not seem to hold for the VaR results. The choice of the window is again very important for the shrinkage method. Unlike for constant correlations, choosing a larger window size does not necessarily mean an improvement of the accuracy of the estimates in case of time-varying correlations. It turns out that a window size of around 50 observations gives the lowest errors and the VaR violation rates are in line with the theory. The results of this study show that applying shrinkage to covariance matrices can indeed improve the estimation considerably both in statistical terms as well as for value-at-risk computations. However, this only holds when the correlations are constant for the entire sample. For time-varying correlations the RiskMetrics^{TM} outperforms the shrinkage method in statistical terms if the true correlations between asset returns are high, but the VaR results are slightly better than those of the shrinkage model only for certain window sizes. RiskMetrics^{TM} renders relatively robust results across different volatility and correlation specifications, but the RiskMetrics^{TM} VaR exceedance rates are relatively high in many cases.

Chapter 6

Summary

Financial markets can be very turbulent and markets change over time. If the changes in dynamic markets are unexpected by investors, they can encounter substantial risk on their investments. Risk on investments can arise from different sources and recent market developments have shown that correlation risk and liquidity risk can have a considerable impact on financial assets. The implications of time-varying correlations and liquidity for financial derivatives has been examined in this thesis.

The impact of jumps, regime switches, and linearly changing correlation term structures on the risk management of basket options has been examined in Chapter 2. First, the results show that there is an asymmetric correlation effect on the value-at-risk of basket options. Second, the time at which a correlation shock occurs during the life of an option is particularly important for hedged basket options. Finally, the square-root-of-time rule can lead to severe underestimation of value-at-risk for basket options with time-varying correlations — for some cases, even by a factor exceeding the minimum regulatory stress factor.

Determining the correlation values between financial assets is a challenge, since correlations are unobservable and have to be estimated by models. Therefore, correlation values are subject to estimation errors. One of the most practically relevant correlation term structures is one containing a sudden jump upward, because sudden increases in asset correlations is often observed during financial crises. Chapter 3 examines the effect of correlation errors for basket options with a sudden increase in the correlations. The results show that the risk measure misestimation effect due to the correlation error can be substantial. An asymmetric risk measure misestimation effect occurs for the unhedged at-the-money and out-of-the-money options, where a negative correlation bias creates a relatively larger deviation from the unbiased risk measures. Moreover, the size and asymmetry effect of the misestimation increases substantially when delta hedging is applied. Finally, the use of the square-root-of-time rule could lead to risk underestimation when correlations are increasing over time.

Liquidity risk of an investment arises if the asset cannot be traded against its model value at a certain time due to, e.g., unfavorable market conditions. Chapter 4 examines the impact of time-varying market liquidity on static barrier option hedging in the foreign exchange market. The results show that the liquidity of currencies depends significantly on market liquidity, but also has an idiosyncratic component. Moreover, time-varying market liquidity characteristics associated with financial crises, such as widening, tightening and jumps of market spreads, have a substantial impact on the liquidity costs of static barrier option hedging. Finally, the magnitude of liquidity costs of the static hedge is highly dependent on the barrier option moneyness.

In Chapter 5 the performance of different models that can be used to estimate correlations has been compared. A Monte Carlo simulation study has been done for the shrinkage techniques applied to covariance matrices as proposed by Ledoit and Wolf (2004) and the results have been compared to the industry standard RiskMetrics^{TM} model. The shrinkage model outperforms RiskMetrics^{TM} in statistical terms as well as in a value-at-risk framework, but only when correlations are constant. For time-varying correlations, the performance of the shrinkage model deteriorates considerably. RiskMetrics^{TM} renders results that are relatively robust to different volatility and correlation specifications.

Nederlandse samenvatting (Dutch Summary)

Financiële markten kunnen erg turbulent zijn. De marktontwikkelingen zijn moeilijk te voorspellen, maar één ding is zeker: de markten zullen veranderen over de tijd. Onverwachte veranderingen in de markt kunnen aanzienlijk risico voor investeringen met zich mee brengen. Recente marktontwikkelingen hebben laten zien dat correlatierisico en liquiditeitsriscio van grote invloed kunnen zijn voor financiële investeringen. In dit proefschrift worden de implicaties van tijdsvariërende marktomstandigheden voor financiële derivaten onderzocht betreffende correlatierisico en liquiditeitsrisico. In de praktijk worden investeringen vaak gedaan in verschillende soorten financiële instrumenten en de correlaties tussen deze instrumenten zijn belangrijk voor het risico management van deze investeringen. Liquiditeitsrisico treedt op wanneer de financiële instrumenten niet verhandeld kunnen worden tegen hun modelwaarde, bijvoorbeeld vanwege ongunstige marktomstandigheden.

In hoofdstuk 2 wordt de impact van verscheidene tijdsvariërende correlatietermijnstructuren op het risico management van basket opties onderzocht. De resultaten laten het volgende zien. Ten eerste, er is een asymmetrisch correlatie-effect op de value-at-risk van de basket opties. Ten tweede, het tijdstip waarop de correlatieschokken optreden gedurende de optie termijn is vooral belangrijk voor opties met een delta hedge. Tenslotte kan de (in de praktijk vaak gebruikte) square-root-of-time rule leiden tot grote onderschatting van value-at-risk voor basket opties met tijdsvariërende correlaties.

Correlaties kunnen in het algemeen niet in de markt geobserveerd worden en moeten geschat worden met behulp van modellen. Het schatten van correlaties kan schattingsfouten veroorzaken. Een van de meest praktisch relevante correlatietermijnstructuren is een waarbij de correlaties via een sprong toenemen, omdat onverwachte toenames in correlaties in de praktijk vaak geobserveerd worden gedurende financiële crises. Hoofdstuk 3 onderzoekt impact van correlatiefouten op basket opties met een correlatiesprong. Uit de resultaten blijkt dat correlatiefouten kunnen leiden tot relatief grote schattingsfouten van de risicomaatstaven. Een asymmetrisch effect treedt op voor de risicomaatstaven van at-the-money en out-of-the-money opties, waarbij een negatieve correlatiefout leidt tot relatief grotere afwijkingen van de risicomaatstaven dan een positieve correlatiefout. Bovendien nemen de grootte en het teken van de relatieve afwijking van de risicomaatstaven als gevolg van een correlatiefout aanzienlijk toe wanneer delta hedging wordt toegepast. Tenslotte wordt aangetoond dat het gebruik van de square-root-of-time rule kan leiden tot onderschatting van risicomaatstaven als correlaties toenemen over de tijd.

Liquiditeitsrisico van een investering ontstaat als de investering niet verhandeld kan worden tegen de modelwaarde op een bepaald tijdstip, bijvoorbeeld vanwege ongunstige marktomstandigheden. Hoofdstuk 4 onderzoekt de impact van tijdsvariërende marktliquiditeit op statisch hedgen van barrière opties in de valutamarkt. De resultaten laten zien dat de liquiditeit van de valuta's significant afhankelijk zijn van de marktliquiditeit, maar ook een idiosyncratisch component bevatten. Bovendien hebben tijdsvariërende marktliquiditeit kenmerken geassocieerd met financiële crises, zoals toenemende, afnemende of sprongen van spreads in de markt, een substantiële impact op de liquiditeitskosten van statisch hedgen van barrière opties. Tenslotte is de grootte van de liquiditeitskosten van een statische hedge sterk afhankelijk van de moneyness van de barrière opties.

In hoofdstuk 5 is een Monte Carlo simulatiestudie uitgevoerd voor het shrinkage covariantiemodel van Ledoit en Wolf (2004) en de resultaten zijn onder meer vergeleken met het RiskMetrics^{TM} model, een model dat in de praktijk veel gebruikt wordt. Het shrinkage model presteert beter dan het RiskMetrics^{TM} model volgens statistische maatstaven en value-at-risk risicomaatstaven, maar alleen als de correlaties constant zijn over de tijd. Voor tijdsvariërende correlaties presteert het shrinkage model aanzienlijk slechter. Het RiskMetrics^{TM} model levert resultaten op die relatief robuust zijn voor verschillende volatiliteit- en correlatiespecificaties.

Bibliography

Acharya, V.V. and L.H. Pedersen, 2005, Asset pricing with liquidity risk, Journal of Financial Economics, 77, 375–410.

Ang, A. and J. Chen, 2002, Asymmetric correlations of equity portfolios, Journal of Financial Economics 63(3), 443–494.

Artzner, P., F. Delbaen, J.M. Eber, and D. Heath, 1999, Coherent measures of risk, Mathematical Finance 9(3), 203–228.

Bank for International Settlements, 2005a, International convergence of capital measurement and capital standards: A revised framework, Technical Report November.

Bank for International Settlements, 2005b, Triennial central bank survey: Foreign exchange and derivatives market activity in 2004, Publication March 2005, Bank for International Settlements, http://www.bis.org.

Bank for International Settlements, 2006, The management of liquidity risk in financial groups, The Working Group review paper May 2006, Bank for International Settlements, http://www.bis.org.

Bauwens, L., S. Laurent, and J.V.K. Rombouts, 2003, Multivariate GARCH models: A survey, Discussion Paper 2003/31, CORE.

Bauwens, L., S. Laurent, and J.V.K. Rombouts, 2006, Multivariate GARCH models: a survey, Journal of Applied Econometrics 21(1), 79–109.

Bennet, M.N. and J.E. Kennedy, 2004, Quanto pricing with copulas, Journal of Derivatives (Fall), 26–45.

Bessembinder, H., 1994, Bid-ask spreads in the interbank foreign exchange markets, Journal of Financial Economics 35(3), 317–348.

Black, F. and M. Scholes, 1973, The pricing of options and corporate liabilities, Journal of Political Economy 81(3), 637–654.

Bollerslev, T., 1986, Generalized autoregressive conditional heteroscedasticity, Journal of Econometrics 31(3), 307–327.

Brenner, M., R. Eldor, and S. Hauser, 2001, The price of options illiquidity, Journal of Finance 56(2), 789–805.

Brigo, D. and F. Mercurio, 2006, *Interest rate models - Theory and practice* (second ed.). Springer.

Brigo, D., F. Mercurio, F. Rapisarda, and R. Scotti, 2004, Approximated moment-matching dynamics for basket-options simulation, Quantitative Finance 4(1), 1–16.

Byström, H.N.E., 2002, Using simulated currency rainbow options to evaluate covariance matrix forecasts, Journal of International Financial Markets, Institutions and Money 12(3), 216–230.

Campbell, A., 2008, A VAR, VAR better thing?, Risk (February).

Campbell, A. and X.L. Chen, 2008, Var counts, Risk (January), 68–71.

Carr, P. and A. Chou, 1997, Breaking barriers, Risk 10(9), 139–145.

Çetin, U., R.A. Jarrow, and P. Protter, 2004, Liquidity risk and arbitrage pricing theory, Finance and Stochastics, 8, 311–341.

Çetin, U., R.A. Jarrow, P. Protter, and M. Warachka, 2006, Pricing options in an extended Black-Scholes economy with illiquidity: Theory and empirical evidence, Review of Financial Studies 19(2), 493–529.

Chordia, T., R. Roll, and A. Subrahmanyam, 2000, Commonality in liquidity, Journal of Financial Economics, 56, 3–28.

Chordia, T., R. Roll, and A. Subrahmanyam, 2001, Market liquidity and trading activity, Journal of Finance 56(2), 501–530.

Connolly, R., C. Stivers, and L. Sun, 2005, Stock market uncertainty and the stock-bond return relation, Journal of Financial and Quantitative Analysis 40(1), 161–194.

Curran, M., 1994, Valuing Asian and portfolio options by conditioning on the geometric mean price, Management Science 40(12), 1705–1711.

Dammers, C.R. and R.N. McCauley, 2006, Basket weaving: the euromarket experience with basket currency bonds, BIS Quarterly Review March, Bank for International Settlements.

Danielsson, J. and J.P. Zigrand, 2005, On time-scaling of risk and the square-root-of-time rule, Technical Report, Working paper, available at www.riskresearch.org.

Deelstra, G., J. Liinev, and M. Vanmaele, 2004, Pricing of arithmetic basket options by conditioning, Insurance, Mathematics and Economics 34(1), 55–77.

Derman, E., D. Ergener, and I. Kani, 1995, Static options replication, Journal of Derivatives (Summer), 78–95.

Diebold, F.X., A. Hickman, A. Inoue, and T. Schuermann, 1997, Converting 1-day volatility to h-day volatiliy: scaling by $\sqrt{h}$ is worse than you think, Technical Report No. 97-34.

Duffie, D. and K.J. Singleton, 2003, *Credit Risk: pricing, measurement, and management* (first ed.), Princeton, NJ, Princeton University Press.

Elton, E.J. and M.J. Gruber, 1973, Estimating the dependence structure of share prices, Journal of Finance, 28, 1203–1232.

Engle, R., 2002, Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models, Journal of Business and Economic Statistics 20(3), 339–350.

Engle, R.F., 1982, Autoregressive conditional heteroscedasticity with estimates of the variance of the U. K. inflation, Econometrica 50(4), 987–1008.

Evans, M.D.D. and R.K. Lyons, 2002, Time-varying liquidity in foreign exchange, Journal of Monetary Economics, 49, 1025–1051.

Eydeland, A. and K. Wolyniec, 2003, *Energy and power risk management* (first ed.), Hoboken, NJ, John Wiley and Sons, Inc.

Fengler, M.R. and P. Schwendner, 2004, Quoting multiasset equity options in the presence of errors from estimating correlations, Journal of Derivatives (Summer), 43–54.

Fleming, M.J., 2003, Measuring treasury market liquidity, Economic Policy Review 9(3), 83–108.

Frey, R. and P. Patie, 2002, Risk management for derivatives in illiquid markets: a simulation study, Working paper, available at SSRN: http://ssrn.com/abstract=300527.

Gibson, M.S. and B.H. Boyer, 1998, Evaluating forecasts of correlation using option pricing, Journal of Derivatives 6(2), 18–38.

Gibson, R. and N. Mougeot, 2004, The pricing of systematic liquidity risk: empirical evidence from the US stock market, Journal of Banking & Finance 28(1), 157–178.

Glasserman, P., 2004, *Monte Carlo methods in financial engineering* (first ed.), New York, Springer-Verlag, Inc.

Glosten, L.R., R. Jagannathan, and D.E. Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, Journal of Finance 48(5), 1779–1801.

Goetzmann, W.N., L. Li, and K.G. Rouwenhorst, 2005, Long-term global market correlations, Journal of Business 78(1), 1–38.

Hansen, P.R. and A. Lunde, 2001, A forecast comparison of volatility models: Does anything beat a GARCH(1,1)?, Working Paper 01-04, Brown University.

Hasbrouck, J., 1999, Security bid/ask dynamics with discreteness and clustering: Simple strategies for modeling and estimation, Journal of Financial Markets 2(1), 1–28.

Hasbrouck, J. and D.J. Seppi, 2001, Common factors in prices, order flows, and liquidity, Journal of Financial Economics, 59, 383–411.

Huberman, G. and D. Halka, 2001, Systematic liquidity, Journal of Financial Research 24(2), 161–178.

Hull, J., 2006, *Options, futures and other derivatives* (sixth ed.), New York, Pearson Prentice Hall.

Jagannathan, R. and T. Ma, 2003, Risk reduction in large portfolios: why imposing the wrong constraints helps, Journal of Finance, 58, 1651–1683.

Jorion, P., 1986, Bayes-stein estimation for portfolio analysis, Journal of Finance, 21, 279–292.

Jorion, P., 2001, *Value at Risk: The new benchmark for managing financial risk* (second ed.). McGraw Hill.

Kupiec, P.H., 1998, Stress testing in a Value at Risk framework, Journal of Derivatives (Fall), 7–24.

Ledoit, O. and M. Wolf, 2003, Improved estimation of the covariance matrix of stock returns with an application to portfolio selection, Journal of Empirical Finance, 10, 603–621.

Ledoit, O. and M. Wolf, 2004, Honey, I shrunk the sample covariance matrix, Journal of Portfolio Management (Summer), 110–119.

Longin, F. and B. Solnik, 1995, Is the correlation in international equity returns constant: 1960-1990?, Journal of International Money and Finance 14(1), 3–26.

Longin, F. and B. Solnik, 2001, Extreme correlation of international equity markets, Journal of Finance 56(2), 649–676.

Lyons, R.K., 2001, *The microstructure approach to exchange rates* (first ed.), Cambridge, Massachusetts, The MIT Press.

Margrabe, W., 1978, The value of an option to exchange one asset for another, Journal of Finance 33(1), 177–186.

McNeil, A.J., R. Frey, and P. Embrechts, 2005, *Quantitative risk management: concepts, techniques, tools* (first ed.), Princeton, NJ, Princeton University Press.

Merton, R.C., 1973, Theory of rational option pricing, Bell Journal of Economics and Management Science 4(1), 141–183.

Milevsky, M.A. and S.E. Posner, 1998, A closed-form approximation for valuing basket options, Journal of Derivatives (Summer), 54–61.

Morgan, J.P., 1996, RiskMetrics—Technical Document, J.P. Morgan Bank, Fourth Edition, New York.

Nalholm, M. and R. Poulsen, 2006a, Static hedging and model risk for barrier options, Journal of Futures Markets 26(5), 449–463.

Nalholm, M. and R. Poulsen, 2006b, Static hedging of barrier options under general asset dynamics: Unification and Application, Journal of Derivatives (Summer), 46–62.

Newey, W.K. and K.D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, Econometrica 55(3), 703–708.

Pástor, Ľuboš. and R.F. Stambaugh, 2003, Liquidity risk and expected stock returns, Journal of Political Economy 111(3), 642–685.

Pellizzari, P., 2005, Static hedging of multivariate derivatives by simulation, European Journal of Operational Research, 166, 507–519.

Pritsker, M.G., 2006, The hidden dangers of historical simulation, Journal of Banking and Finance 30(2), 561–582.

Skintzi, V.D., G. Skiadopoulos, and A.P.N. Refenes, 2005, The effect of misestimating correlation on Value at Risk, Journal of Alternative Investments (Spring), 66–82.

Stein, C., 1955, Inadmissibility of the usual estimator for the mean of a multivariate normal distribution, in Proceedings of the third Berkeley symposium on mathematical statistics and probability, Vol. IV, (Berkeley, University of California Press), 2111–2245.

Taleb, N., 1997, *Dynamic Hedging: managing vanilla and exotic options* (First ed.), New York, John Wiley & Sons.

Wong, A.S.K., 2006, Basel II and the risk management of basket options with time-varying correlations, International Journal of Central Banking 2(4), 1–32.

Zakoian, J.M., 1994, Threshold heteroskedastic models, Journal of Economic Dynamics and Control 18(5), 931–955.

Zhang, P., 1998, *Exotic options: a guide to second generation options* (second ed.). World Scientific.

The Tinbergen Institute is the Institute for Economic Research, which was founded in 1987 by the Faculties of Economics and Econometrics of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam and Vrije Universiteit Amsterdam. The Institute is named after the late Professor Jan Tinbergen, Dutch Nobel Prize laureate in economics in 1969. The Tinbergen Institute is located in Amsterdam and Rotterdam. The following books recently appeared in the Tinbergen Institute Research Series:

395 . A. LLENA NOZAL, *On the dynamics of health, work and socioeconomic status.*

396 . P.D.E. DINDO, *Bounded rationality and heterogeneity in economic dynamic models.*

397 . D.F. SCHRAGER, *Essays on asset liability modeling.*

398 . R. HUANG, *Three essays on the effects of banking regulations.*

399 . C.M. VAN MOURIK, *Globalisation and the role of financial accounting information in Japan.*

400 . S.M.S.N. MAXIMIANO, *Essays in organizational economics.*

401 . W. JANSSENS, *Social capital and cooperation: An impact evaluation of a womens empowerment programme in rural India.*

402 . J. VAN DER SLUIS, *Successful entrepreneurship and human capital.*

403 . S. DOMINGUEZ MARTINEZ, *Decision making with asymmetric information.*

404 . H. SUNARTO, *Understanding the role of bank relationships, relationship marketing, and organizational learning in the performance of peoples credit bank.*

405 . M.. DOS REIS PORTELA, *Four essays on education, growth and labour economics.*

406 . S.S. FICCO, *Essays on imperfect information-processing in economics.*

407 . P.J.P.M. VERSIJP, *Advances in the use of stochastic dominance in asset pricing.*

408 . M.R. WILDENBEEST, *Consumer search and oligopolistic pricing: A theoretical and empirical inquiry.*

409 . E. GUSTAFSSON-WRIGHT, *Baring the threads: Social capital, vulnerability and the well-being of children in Guatemala.*

410 . S. YERGOU-WORKU, *Marriage markets and fertility in South Africa with comparisons to Britain and Sweden.*

411 . J.F. SLIJKERMAN, *Financial stability in the EU.*

412 . W.A. VAN DEN BERG, *Private equity acquisitions.*

413 . Y. CHENG, *Selected topics on nonparametric conditional quantiles and risk theory.*

414 . M. DE POOTER, *Modeling and forecasting stock return volatility and the term structure of interest rates.*

415 . F. RAVAZZOLO, *Forecasting financial time series using model averaging.*

416 . M.J.E. KABKI, *Transnationalism, local development and social security: the functioning of support networks in rural Ghana.*

417 . M. POPLAWSKI RIBEIRO, *Fiscal policy under rules and restrictions.*

418 . S.W. BISSESSUR, *Earnings, quality and earnings management: the role of accounting accruals.*

419 . L. RATNOVSKI, *A Random Walk Down the Lombard Street: Essays on Banking.*

420 . R.P. NICOLAI, *Maintenance models for systems subject to measurable deterioration.*

421 . R.K. ANDADARI, *Local clusters in global value chains, a case study of wood furniture clusters in Central Java (Indonesia).*

422 . V.KARTSEVA, *Designing Controls for Network Organizations: A Value-Based Approach.*

423 . J. ARTS, *Essays on New Product Adoption and Diffusion.*

424 . A. BABUS, *Essays on Networks: Theory and Applications.*

425 . M. VAN DER VOORT, *Modelling Credit Derivatives.*

426 . G. GARITA, *Financial Market Liberalization and Economic Growth.*

427 . E.BEKKERS, *Essays on Firm Heterogeneity and Quality in International Trade.*

428 . H.LEAHU, *Measure-Valued Differentiation for Finite Products of Measures: Theory and Applications.*

429 . G. BALTUSSEN, *New Insights into Behavioral Finance.*

430 . W. VERMEULEN, *Essays on Housing Supply, Land Use Regulation and Regional Labour Markets.*

431 . I.S. BUHAI, *Essays on Labour Markets: Worker-Firm Dynamics, Occupational Segregation and Workplace Conditions.*

432 . C. ZHOU, *On Extreme Value Statistics.*

433 . M. VAN DER WEL, *Riskfree Rate Dynamics: Information, Trading, and State Space Modeling.*

434 . S.M.W. PHLIPPEN, *Come Close and Co-Create: Proximities in pharmaceutical innovation networks.*

435 . A.V.P.B. MONTEIRO, *The Dynamics of Corporate Credit Risk: An Intensity-based Econometric Analysis.*

436 . S.T. TRAUTMANN, *Uncertainty in Individual and Social Decisions: Theory and Experiments.*

437 . R. LORD, *Efficient pricing algorithms for exotic derivatives.*

438 . R.P. WOLTHOFF, *Essays on Simultaneous Search Equilibrium.*

439 . Y.-Y. TSENG, *Valuation of travel time reliability in passenger transport.*

440 . M.C. NON, *Essays on Consumer Search and Interlocking Directorates.*

441 . M. DE HAAN, *Family Background and Children's Schooling Outcomes.*

442 . T. ZAVADIL, *Dynamic Econometric Analysis of Insurance Markets with Imperfect Information*

443 . I.A. MAZZA, *Essays on endogenous economic policy*

444 . R. HAIJEMA, *Solving large structured Markov Decision Problems for perishable-inventory management and traffic control*