



**Norges
Handelshøyskole**

*Norwegian School of Economics
and Business Administration*

**FOUR ESSAYS ON SECURITY DESIGN:
MOTIVATION AND IMPLEMENTATION**

By
Tarres G. Trovik

**SUBMITTED FOR THE DEGREE OF
DR. OECON
AT
NORWEGIAN SCHOOL OF
ECONOMICS AND BUSINESS ADMINISTRATION
BERGEN, NORWAY
AUGUST 2001**

ISBN 82-405-0062-5

Table of Contents

Table of Contents	iii
Acknowledgements	v
Introduction	1
1 Optimal Hedging of Contingent Exposure: The Importance of a Risk Premium	7
1.1 Introduction	7
1.2 Economic Model	10
1.2.1 The Initial Problem	10
1.2.2 The Hedging Problem	12
1.3 Preference and Belief Specifications	13
1.3.1 The Optimal Initial Payoff Under Preference Restrictions	13
1.3.2 Belief Specification	15
1.4 The Optimal Hedging Claim	16
1.4.1 Quadratic Utility	16
1.4.2 Exponential Utility	18
1.4.3 A Note on the Classic Hedging Analysis with Exponential Utility	19
1.4.4 Logarithmic Utility	20
1.5 Concluding Remarks	24
2 Optimal Multivariate Exposure: Hedging Currency Risk of Foreign Equity Investments	25
2.1 Introduction	25
2.2 Optimal Exposure	27
2.3 Optimal Multivariate Payoff when Prices are Lognormally Distributed	29
2.3.1 Joint Distribution of Assets and Optimal Wealth	31
2.3.2 Diversification	33
2.4 Hedging Currency Risk of Foreign Equity Exposure	34
2.4.1 A Comment on Siegel's Paradox	38
2.5 Concluding Remarks	40

3	Efficient Replication of Derivatives Defined on Portfolios	41
3.1	Introduction	41
3.2	The Financial Market	43
3.3	The Problem	44
3.3.1	A Standard Approximation	46
3.3.2	The Suggested Model	47
3.3.3	The Geometrically Weighted Portfolio	54
3.4	Numerical Testing	58
3.5	Concluding Remarks	65
4	Options on Active Portfolios: Pricing Under Semi-Lognormality	67
4.1	Introduction	67
4.2	Financial Model	70
4.3	Standard Approximation of European Portfolio Exchange Options	70
4.4	Externalizing the Trigger Assets	72
4.4.1	Pricing a Four Asset Exchange Option	74
4.4.2	Assuming a Joint Lognormal Distribution	78
4.5	Pricing Under a Joint Lognormal-Normal Distribution	80
4.5.1	Pricing a Portfolio Spread Option	85
4.6	Numerical Testing	86
4.7	Concluding Remarks	89
4.8	Appendix: Parameters of the Seven Cases	90
	Bibliography	92

Acknowledgements

This work could not have been completed without the support and help from a number of people.

First comes my beloved Tina; without her steady encouragement and never failing optimism this project surely would have ended in default.

I am furthermore indebted to my supervisor Professor Gunnar Stensland who's well placed questions, suggestions and hints have contributed greatly to the present work. I also learned a lot from Svein-Arne Persson and gained from his encouragement.

The Department of Finance and Management Science at NHH has been a very stimulating place to conduct this research. The stimulating teachers, students and friends in my period at NHH are too numerous to be mentioned in full. Special thanks to Thore Johnsen who woke my interest in the theory of finance. Other main contributors in keeping that interest alive includes Knut Aase, Bernt Øksendal, Steinar Ekern and Terje Lensberg.

A number of friends and colleagues have contributed through little comments that were far more important than they probably realized, in particular Hans Olav Husum, Anders Buvik, Jørgen Haug, Eirik Gaard Kristiansen, Bernt Christian Brun and finally Bernt Arne Ødegaard who also spent a lot of time leading me deeper into the fascinating world of programming.

The bulk of this work was performed with a grant from The Norwegian School of Economics and Business Administration during the period from June 1994 to April 1997, and with a grant from my employer Norges Bank Investment Management, from October 1999 to June 2000. The latter made possible with the generous cooperation from my boss Morten Jensen. University of California at Berkeley kindly accepted me as visitor in 1995/1996.

Finally, I would like to express my gratitude to my family; your support continues to be important to me.

Bergen, August 2001
Tørres G. Trovik

Introduction

The questions of *Why* and *How* continue to keep scientists busy. In the narrow scope of this thesis, the questions translate to: Why does an investor want to obtain a certain distribution of future payoffs? How can a certain distribution of payoffs be priced and obtained, given that it is desired by someone? The answers to these questions are obviously intimately related, e.g., knowing "why" will in a market equilibrium setting provide the answer to "how", and vice versa. It is an example of a primal and a dual problem.

Within the theory of finance, and in particular continuous-time finance, one can argue that more attention has been paid to the latter question, i.e., the implementation issue. Following the seminal papers by Samuelson (1965), Black and Scholes (1973), Merton (1973), Harrison and Kreps (1979) and Harrison and Pliska (1981) who all are main contributors to the modern theory of asset pricing, a substantial amount of research has been devoted to applying that framework under increasingly relaxed assumptions and to increasingly exotic distributions of payoffs.

Less attention has been paid to the motivation for investing in exotic, or even standard, derivatives. Notable exceptions, however, include Brennan and Solanki (1981) and Leland (1980). In the study by Brennan and Solanki (1981) the emphasis is on what kind of insurance contract, i.e., a non-linear payoff profile, an investor should buy, given his preferences and beliefs. Their results are based on the assumptions of a "risk-neutral valuation relationship" and lognormality in the return distribution. The underlying portfolio is arbitrary and *market* preferences and beliefs are not introduced. This contrasts with the analysis in Leland (1980) where optimal risk sharing contracts between an individual investor and the hypothetical *representative investor* who supports market prices are described.

Moreover, the seminal paper by Cox and Huang (1989) demonstrated the close link between the dynamic portfolio choice problem as studied by Merton (1971) and others, and a static optimization of a future payoff profile. They showed that the dynamic trading strategy

replicating the optimal payoff is equivalent to the solution of a corresponding dynamic optimization problem. Thus, in this sense the vast literature on portfolio theory can be seen as a study of optimal, non-linear payoff profiles¹.

Demand for non-linear payoff profiles is not sufficient motivation for investing in derivatives by itself. In the standard asset pricing framework markets are dynamically complete by assumption. This in effect makes derivatives redundant securities. Thus, the demand for a non-linear payoff profile through the purchase of a derivative must be explained by e.g. a convenience factor which is external to our models.

The first two chapters of this thesis deal with the question of *why*. In a standard utility-based setting we develop utility-maximizing payoff profiles or derivatives, and discuss how these payoff profiles are affected by preferences and beliefs. The setting is similar to that of Brennan and Solanki (1981), but the discussion is confined to two different situations. In the first chapter we analyze optimal payoff profiles for a hedging problem where the exposure to be hedged is uncertain itself. In the second chapter the approach is generalized to apply to more than one risky asset. Our example involves the hedging of a foreign equity investment. Thus, in this setting the final wealth of the agent depends multiplicatively on the future prices of two risky assets.

The third and fourth chapter address *how* questions. Here, the demand for the payoff profiles analyzed is not an issue. Rather, well known profiles are chosen, but the underlying "security" is a portfolio and not a single asset. This represents a challenge in a lognormal world, because the probability distribution of sums of lognormal variates is unknown. In the third chapter we look at a perpetual put option on a portfolio and propose and test a replicating trading scheme. In the last chapter the focus is on European spread options on the difference between two portfolios, and we propose and test an approximation formula for its fair value.

Following this attempt to provide an overall framework in which to place the different essays, the remainder of this introduction summarizes the contents of each chapter.

The focus of the first chapter is an agent's choice of a payoff profile when hedging an underlying exposure. We look at a situation where the exposure to be hedged is contingent on a non-marketed event. An example is a contractor who acquires a currency exposure if he wins a bidding contest for a foreign project. In the literature, there are many examples where the

¹A payoff profile is a function relating the prices of underlying securities to the payoff received by the holder of a contract or replicating strategy. Examples include $f(S) = S$ which is the linear payoff profile of constantly holding the underlying security, and $f(S) = \max(K - S, 0)$ which is the non-linear payoff of a call option. f is assumed to be unknown.

optimal use of particular hedging instruments, such as forwards and options, is analyzed. In our work, however, the optimal payoff profile is endogenous. Thus, the analysis is not confined to the optimal use of one particular hedging instrument; rather, the optimal payoff is derived on the basis of the agent's preferences and beliefs. As a special case we show that when the exposure is *not* contingent, i.e., a standard exposure, the standard approach where the analysis is confined to one particular hedging instrument produce a misleading answer.

We show, for various preferences, how the size of a risk premium, i.e., the agent's subjective beliefs of expected return versus the risk-free interest rate, affects the degree of non-linearity in the optimal hedging instrument. This result is in contrast to known results for contingent exposure in the case of a zero risk premium. We interpret the optimal contingent claim as a sum of an initial component and a hedge component. The initial component is the claim the investor would choose in order to maximize his utility, irrespective of the need to hedge. The hedge component is the claim the investor uses to hedge the exposure from the bidding contest. In the literature these two components are sometimes termed "speculation demand" and "hedging demand".

In the case of exponential utility we demonstrate that the hedge component is independent of the risk premium. In case of quadratic and logarithmic utility, the hedge component exhibits similar non-linearities to the initial component. We may interpret the pure hedge under both quadratic and logarithmic utility as forward contracts plus the present value of the forward price invested in the same manner as the initial component.

In the second chapter, we generalize both the economic model and the approach used in the first chapter, to allow for more than one risky asset. Our aim is to characterize optimal exposure to two risky assets.

In a standard mean variance approach, the goal is to find optimal holdings of each asset, implicitly assuming that a long or short holding of the (underlying) asset is the only option. Our approach, on the other hand, is to analyze the optimal exposure, described as a function of the two risky asset prices. Thus, the payoff profile in this case is a surface in three-dimensional space.

In the first sections of the chapter we study optimal straightforward exposure, or what we called the initial component in chapter 1. The analysis is similar to that of Brennan and Solanki (1981), but they only analyze a one-asset economy. Our explicit formulation of an optimal bi-variate claim lends itself to illustrations of diversification as well as the probability distribution of optimal wealth. We illustrate that rather than to exogenously assume both

a joint distribution of prices and aggregate wealth, as well as preferences of a representative agent, the distribution of aggregate wealth follows endogenously from assumptions on the distribution of prices and preferences of a representative agent. Cox and Huang (1989) also analyze a multi-asset economy, but they do not characterize the optimal payoff function.

In section four we introduce a particular hedging problem and use our approach much in the same way as in the first chapter. However, here our case is that of an investor buying equity abroad, thus his exposure is the product of the stock price and the currency exchange rate. The exposure is not contingent, as it was in chapter one. The investor has different expectations regarding risk and return for the currency and the stock, so he should not necessarily hold the same exposure to the two assets. We analyze how the utility-maximizing hedge should be structured, given preferences and beliefs, and demonstrate conditions for when a claim of "quanto" type is demanded. A "quanto" claim is a derivative that pays off a foreign price in home currency units. For example, a Norwegian quanto futures contract on the American S&P 500 index pays off NOK 1417 when the futures price is 1417.

This chapter also contains a comment on the Siegel (1972) Paradox, and explains why it does not affect our optimal hedging results, as it does in Black (1990). The analysis relies on the change of numeraire theorem introduced by Geman, El Karoui, and Rochet (1995).

In the third chapter we study efficient or cost-effective ways of replicating a perpetual put option on a portfolio with constant holdings. In this chapter we are not rigorously analyzing why an investor demands such an option, but treat the need as given. However, one intuitive motivation can be found in a risk-management setting. Suppose an active manager put together a portfolio of (long only or short only) positions that he believes will perform well, but he is not able to define any explicit time-horizon for his beliefs. He wants to avoid downside risk, so a perpetual put option on the portfolio is one type of stop-loss strategy².

When designing a replicating strategy for an option, it is necessary to find the sensitivity of the value of the option to a change in the value of the underlying asset. This is called the "delta" of the option. In our case, the underlying asset is a portfolio consisting of several (correlated) assets. Thus, there is a delta for each asset, i.e., the derivative of the option value with respect to each asset price.

Because the sum of lognormally distributed prices has an unknown probability distribution, the underlying asset in our case has unknown distribution. A common approximation (see

² Admittedly a more realistic example would be to let the horizon be indefinite for a much smaller time interval, say 14 days to 12 months, than from now to infinity, which is the situation represented with a perpetual option. Still, one often finds that practitioners make financial investments without any explicit time horizon.

Hull (2000)), is to pretend that the portfolio is lognormally distributed, as if the portfolio was a single asset. To obtain the deltas from such an approximated value, however, is not without its problems. With a straightforward calculation, the delta with respect to each asset is the derivative with respect to the portfolio, times the derivative of the portfolio with respect to the asset-value. The latter derivative is one, because the portfolio is the sum of the values of the constituent assets; thus the delta with respect to a single asset is the same across all assets in the portfolio.

This is not intuitively appealing for the following reason: Even if constant volatility is assumed for each asset, the volatility of the portfolio will change over time, for example when the market movement of a high-volatility asset causes a change in that asset's portfolio weight. And we know that the volatility of the underlying asset affects the value of the delta. Thus if the portfolio consists of assets with different volatilities, it is reasonable to suppose that the delta for each asset will be different across assets.

Another common way of approximating derivatives on portfolios is to approximate the value of the portfolio by means of a geometrically weighted average, which we know is lognormally distributed (see e.g. Zhang (1997)). When the deltas are calculated using this approach, they are found to be different for each asset³, but it is still only the portfolio volatility and not the volatility of the individual asset, that is involved in the expression for the deltas.

In the chapter we suggest a new method of computing the delta for each asset in the portfolio that explicitly takes into account the parameters of the process for that asset. Then we test the efficiency of a replicating strategy based on this approach, compared with the geometric average approach outlined above.

The test is performed with Monte Carlo simulation. The proposed strategy compares favorably by displaying a greatly decreased turnover, thus reducing the cost of the option in a world with transaction costs. Our strategy yields a similar, or slightly better, ability to obtain the payoff of the option.

In the fourth and last chapter, the problem has the same source as that in chapter three: The sum of lognormal prices has an unknown distribution. However, in this chapter we study a different kind of option. Moreover, replication is not the issue; rather, we propose and test a new way of approximating its value.

The option we study is a portfolio spread option, that is, an option, European call or put,

³The difference is here due to the derivative of a geometric average with respect to one of its constituent values.

on the difference between two portfolios. Other ways of describing the option are as a call option on a portfolio with another portfolio as strike, or as an option on a long-short portfolio.

This kind of option could be of particular interest to an active portfolio manager. Normally an active manager is told to beat a benchmark portfolio defined by his sponsors by applying his predictive skills. A portfolio spread option is thus of interest when it is relevant to apply risk management relative to a benchmark portfolio.

An exact pricing formula is not available in a Black and Scholes world because the distribution of sums of lognormal variables is unknown. The standard way of approximating its value is again by pretending the portfolios are lognormally distributed.

In this chapter we propose a new approach whereby we first decompose the option into a sum of options using a generalized version of a technique used in chapter three. This decomposition allows us to assume different probability distributions for individual assets and portfolios. Then we use one particular distribution, called a joint lognormal-normal distribution or a semi-lognormal distribution, to calculate the approximated value. This distribution was described in Crow and Shimizu (1988), but to my knowledge was not applied in finance until it was used by Camara and Stapleton (1998).

We test our approach by comparing our approximated option values for different parameters with the "true value" obtained by Monte Carlo simulation and with the standard approximation approach. The results are not very encouraging, as our approximated values are very similar to the values obtained using the standard approach. However, our approach makes it possible to approximate types of options very easily, which is not possible with the standard approach. Pricing formulas for several rather exotic derivatives are developed as preliminary results.

The four chapters have been written as self-contained articles. They may be read separately, in any order favored by the reader.

Chapter 1

Optimal Hedging of Contingent Exposure: The Importance of a Risk Premium

Co-author: Svein-Arne Persson¹

The focus of this chapter is how a non-zero risk premium affects an economic agent's optimal hedging decision when exposed to a non-marketed event. The analysis is not confined to the optimal use of one particular hedging instrument; rather, the optimal payoff is derived on the basis of the agent's preferences. We show, for various preferences, how the size of a risk premium affects the degree of non-linearity in the optimal hedging instrument. This result is in contrast to known results for contingent exposure in the case of a zero risk premium. We demonstrate an inefficacy of the approach of confining the analysis to one particular hedging instrument in the case of standard exposure.²

1.1 Introduction

Faced with a situation of financial risk, an agent may or may not choose to engage in hedging activities. An agent exposed to financial risk with probability one is said to be faced with a *standard exposure*. In this chapter we study an exposure contingent on a non-marketed event, as illustrated by the following standard example.

Consider a contractor who bids 1 unit of foreign currency in an auction for a construction project abroad. The probability of winning the contract is assumed to be p . Hence, with probability p the contractor will receive a fixed amount of foreign currency at a fixed future date, say time 1, and with probability $(1 - p)$ he will receive nothing. Currency and derivatives

¹Norwegian School of Economics and Business Administration, Institute of Finance and Management Science, Helleveien 30, N-5035 Bergen-Sandviken - Norway. E-mail: Svein-Arne.Persson@nhh.no

²Published in The Journal of Futures Markets, Vol. 20, No. 9, 823-841(2000)

of currency are by assumption traded in a competitive financial market. However, no assets contingent on the outcome of the auction are available for trade. The auction thus represents a non-marketed event. We assume that the probability p is independent of the market value of the currency.

Note that by letting p be equal to one or zero, we recover a setting of standard exposure or no exposure, respectively. The purpose of this chapter is to investigate how the contractor should design an optimal hedge. A complete hedge, i.e., a hedge eliminating all uncertainty, is not possible for p strictly positive and less than 1 because then the claim is dependent on a non-marketed event.

In the hedging literature, a frequently used approach is to confine the analysis to the use of one particular hedging instrument. That is, a decision maker first chooses an instrument like a forward, futures, or option contract; the optimal use of that instrument in order to manage the financial risk is then analyzed.

In this chapter we do not take the hedging instrument for given, but derive the optimal hedging instrument characterized by its payoff as a function of the value of the underlying asset. We focus on how this optimal payoff is affected by the presence of a non-marketed exposure and a non-zero risk premium.

The present work is related to that reported by Steil (1993). He works within the contingent exposure setting, and his analysis is not confined to one particular financial derivative. However, he assumes unbiased prices of contingent claims. Several models of optimal hedging (e.g. Lapan, Moschini, and Hanson (1991) and Moschini and Lapan (1995)) are based on the assumption of unbiased prices of contingent claims in the sense that a futures price or option premium equals the agent's expected payoff from the respective contracts. A risk premium is then defined in terms of deviations between the market price and the expected payoffs, defined as biases in the cited papers.

We, on the other hand, introduce the risk premium directly on the underlying assets as in a standard Black and Scholes economy. Here redundant contingent claims can be priced by arbitrage, thus any risk premium on the contingent claim is a function of the risk premium of the underlying asset. By introducing a non-zero risk premium we are able to study the combined effects of a risk premium and a non-marketed exposure on the shape of the optimal payoff function.

This chapter is also related to works on optimal portfolio insurance by Leland (1980) and Brennan and Solanki (1981). Leland (1980) characterizes investors in terms of how their risk

aversion and expectations differ from the market average, and focuses on who will benefit from non-linear payoffs. In particular, he analyzes the demand for portfolio insurance. Brennan and Solanki (1981) provide the explicit form of this optimal payoff for agents with preferences exhibiting hyperbolic absolute risk aversion (HARA). This problem is further analyzed by Carr and Madan (1997) using a similar framework as in the present chapter. Relating these papers to our work, their setting resembles the case of no exposure, i.e., our case of $p = 0$.

Other related works within the hedging literature include Feiger and Jacquillat (1979), Briys, Crouhy, and Schlesinger (1993), Moschini and Lapan (1992), Eaker and Grant (1985), Giddy (1985), Lehrbass (1994), Stulz (1984), Kerkvliet and Moffet (1991), and Adler and Detemple (1988).

Our approach yields three important insights. First, we confirm the well known result that a non-zero risk premium induces an expected utility maximizer to hold currency independently of the project considered. Thus, it is natural to consider two components of an optimal contingent claim. The *initial component* is the optimal position prior to the considered auction. The *optimal hedge component* is the contingent claim required to optimally alter the exposure from the project.

Second, we show that both the presence of a non-marketed exposure as well as the size of the risk premium has a direct effect on the shape of the optimal payoff. We analyze optimal payoffs for the cases of quadratic, exponential and logarithmic utilities. For $p = 1$, i.e., the case of standard exposure, the optimal hedge has a linear payoff only for the case of exponential utility. When the exposure is contingent, i.e., p different from 1 or 0, the size of p affects the curvature of the optimal payoff for all preferences considered. The initial component is non-linear for all analyzed preferences and is, of course, independent of p . When the risk premium approaches zero, the optimal hedge becomes increasingly linear. When the risk premium is equal to zero, a linear hedge is optimal and the initial component is zero.

Third, we illustrate limitations of the standard approach, where the analysis is confined to the optimal use of particular contingent claims, like forwards, futures or options. The standard procedure is to find an optimal fraction of the exposure to hedge with one or several favored instruments, and normally these fractions vary with the size of the risk premium. We show that a change in the risk premium demands a different payoff rather than a different exposure to the same payoff. A replicating strategy for the optimal claim can be implemented either dynamically with the underlying assets, or by a static holding of marketed contingent claims (e.g., Carr and Madan (1997)). Rather than choosing the instruments and their properties

somewhat arbitrarily prior to the optimization, knowledge of the optimal payoff guides us to choose the replicating instruments.

The chapter is organized as follows: Section 1.2 describes the economic model. In Section 1.3 preferences and belief specification are outlined. In Section 1.4 the optimal hedging claims are analyzed under different preferences of the HARA type. Section 1.5 contains some concluding remarks.

1.2 Economic Model

The object of study is the optimal design of a contingent claim based on a contractor's preferences and beliefs. Our set-up differs from the standard Black and Scholes economy by allowing a non-marketed outcome. The mentioned contractor receiving a payment in foreign currency with some positive probability is our standard example of this friction.

The existence of a positive risk premium in the financial market encourages the contractor to undertake some investments independently of the project considered. In the following subsection relevant results for the contractor's *initial problem* by Brennan and Solanki (1981), Carr and Madan (1997), and Cox and Huang (1989) are collected.

1.2.1 The Initial Problem

The contractor makes his investments at the initial date, time 0, and all uncertainty is resolved at the terminal date, time 1. The contractor is not allowed to rebalance his portfolio between time 0 and time 1. The contractor's preferences are represented by a von Neuman-Morgenstern utility function ($U' > 0$ and $U'' < 0$) for time 1 wealth only.

The only source of uncertainty is the foreign exchange rate at time 1, denoted by S (number of domestic units per foreign unit). Formally, S is a random variable with the non-negative real numbers as support. The initial exchange rate is given by the constant S_0 .

The contractor seeks the optimal allocation of terminal wealth by investing in a claim contingent on the foreign exchange rate. This claim is denoted by $g(S)$. In subsection 2.2 this initial claim needs to be supplemented by another claim, termed a hedging claim, because the contractor is faced with an additional currency exposure from the described project.

The contractor's initial wealth is denoted by W_0 and the domestic interest rate r_d is assumed constant.

Initially, the contractor wants to invest in the contingent claim solving the problem

$$\max_{g(s)} E[U(W)] = \max_{g(s)} \int_0^{\infty} U(g(s))f(s)ds, \quad (1.2.1)$$

i.e, he maximizes expected utility on the basis of his beliefs about the exchange rate at time 1, represented by the probability density $f(\cdot)$. He is constrained by his time zero wealth W_0 and the market prices prevailing at time zero,

$$\int_0^{\infty} e^{-\tau_d} g(s)q(s)ds = W_0. \quad (1.2.2)$$

From standard theory of finance we know that absence of arbitrage opportunities and dynamic completeness is sufficient within our model for the existence of a unique risk-adjusted probability density $q(\cdot)$.

The Lagrangean of the problem is

$$\mathcal{L} = \int_0^{\infty} U(g(s))f(s)ds - \lambda \left(\int_0^{\infty} e^{-\tau_d} g(s)q(s)ds - W_0 \right),$$

where λ is the constant Lagrange-multiplier. The first order condition of this problem is

$$\lambda = U'(g(S)) \frac{f(S)}{q(S)e^{-\tau_d}}. \quad (1.2.3)$$

Consider the Arrow-Debreu security represented by the indicator function $1\{\omega \in [s, s+ds)\}$, where ω represents the state of the world. This security pays one unit if and only if a given state ω occurs. Its expected payoff is $f(s)ds$, whereas the time zero market price is $q(s)e^{-\tau_d}ds$. From expression (1.2.3) we see that the agent chooses the payoff such that the expected gross³ return of all Arrow-Debreu securities, multiplied by the marginal utility of the optimal claim, is constant through different states. This is an established result in economics that can be traced back to Borch (1962).

The optimal payoff $g(S)$ in terms of the inverse marginal utility function can be determined from expression (1.2.3) as

$$g(S) = [U']^{-1} \left(\lambda \frac{q(S)e^{-\tau_d}}{f(S)} \right). \quad (1.2.4)$$

The claim $g(S)$ represents the payoff of the investments the contractor would like to undertake whether or not he is exposed to additional foreign exchange risk. In the following sub-section the described hedging problem is introduced.

³Recall that $\frac{\text{payoff}}{\text{price}} = \text{gross return} = 1 + \text{net return}$.

1.2.2 The Hedging Problem

Now the contractor's total exposure when faced with the prospect of receiving one unit of S with probability p is considered. In addition to the initial claim $g(S)$ from expression (1.2.4), the contractor now also invests in a hedging claim, denoted by $h(S)$, to account for the additional currency exposure. The total optimal payoff $y(S)$, excluding the potential currency payoff from the project, is then the sum of the initial claim and the hedging claim, i.e. $y(S) = h(S) + g(S)$.

The contractor now seeks the contingent claim that solves the following problem

$$\max_{y(s)} \mathbb{E}[U(W)] = \max_{y(s)} \int_0^{\infty} [pU(s + y(s)) + (1 - p)U(y(s))] f(s) ds, \quad (1.2.5)$$

subject to the budget constraint

$$\int_0^{\infty} e^{-ra} y(s) q(s) ds = W_0. \quad (1.2.6)$$

Note in the formulation of the problem (1.2.5) and (1.2.6), that if $p = 0$, i.e., the probability of receiving S is zero, $y(S)$ is equal to $g(S)$ given by (1.2.1). Note also that if $p = 1$, the problem is reduced to an ordinary hedging problem where the exposure does not depend on a non-marketed event.

We form the Lagrangean,

$$\mathcal{L} = \int_0^{\infty} [pU(s + y(s)) + (1 - p)U(y(s))] f(s) ds - \lambda \left(\int_0^{\infty} e^{-ra} y(s) q(s) ds - W_0 \right),$$

where again λ represents the Lagrange-multiplier. The first order condition for this problem is

$$pU'(S + y(S)) + (1 - p)U'(y(S)) = \lambda \frac{q(S)e^{-ra}}{f(S)}. \quad (1.2.7)$$

The structure of this equation is somewhat more complex than the corresponding equation (1.2.3) for the case without the additional non-marketed exposure.

The left-hand side of (1.2.7) can be interpreted as the expected marginal utility of the optimal claim, including expectations with respect to non-marketed outcomes. Thus, it has a close resemblance to the analysis of background risk (e.g., Franke, Stapleton, and Subrahmanyam (1998)).

1.3 Preference and Belief Specifications

In this section preferences are specified. We apply a general HARA utility function with various further restrictions on the parameters to, for example, obtain results for utility functions with constant absolute and constant relative risk aversion.

Table 1.1: Utility functions.

Name	General HARA	Quadratic	Exponential	Logarithmic
Param.		$\tau = 2b$ $\gamma = -1$ $A = -1$	$\gamma \uparrow \infty$ $A = -1$	$\tau = 1$ $\gamma \rightarrow 1$ $A = 1$
$U(W)$	$\frac{\tau}{1-\gamma} (\frac{\tau}{\gamma}W - A)^{1-\gamma}$	$W - bW^2$	$-e^{-\tau W}$	$\ln(W)$
$U'(W)$	$\tau (\frac{\tau}{\gamma}W - A)^{-\gamma}$	$1 - 2bW$	$\tau e^{-\tau W}$	$\frac{1}{W}$
$U''(W)$	$-\tau^2 (\frac{\tau}{\gamma}W - A)^{-\gamma-1}$	$-2b$	$-\tau^2 e^{-\tau W}$	$-\frac{1}{W^2}$
Rest.	$\frac{\tau}{\gamma}W \geq A$	$b > 0, W \leq \frac{1}{2b}$	$\tau > 0$	
ARA	$\frac{\tau}{\frac{\tau}{\gamma}W + A}$	$\frac{2b}{1-2bW}$	τ	$\frac{1}{W}$
RRA	$\frac{\tau W}{\frac{\tau}{\gamma}W + A}$	$\frac{2bW}{1-2bW}$	τW	1

1.3.1 The Optimal Initial Payoff Under Preference Restrictions

Table 1 displays most important characteristics of the preference specifications considered in this study.

To simplify the presentation of the solution of equation (1.2.4), the following security payoffs are defined, each corresponding to the four different utility functions of Table 1:

$$c_H(S) = \left(\frac{f(S)}{q(S)e^{-\tau d}} \right)^{\frac{1}{\gamma}},$$

$$c_Q(S) = \left(\frac{f(S)}{q(S)e^{-\tau d}} \right)^{-1},$$

$$c_E(S) = \ln \left(\frac{f(S)}{q(S)e^{-\tau d}} \right),$$

and

$$c_L(S) = \frac{f(S)}{q(S)e^{-\tau d}}.$$

The corresponding time zero market prices are $\pi_H = \int_0^\infty c_H(s)e^{-r_d q(s)}ds$, $\pi_Q = \int_0^\infty c_Q(s)e^{-r_d q(s)}ds$, $\pi_E = \int_0^\infty c_E(s)e^{-r_d q(s)}ds$, and $\pi_L = \int_0^\infty c_L(s)e^{-r_d q(s)}ds = 1$.

Here $c_H(S)$, $c_Q(S)$, and $c_L(S)$ are non-negative in all states ($f(\cdot)$ is a probability density) and represent the expected gross return of an Arrow-Debreu security raised to a power depending on the agent's risk aversion. In contrast, the payoff $c_E(S)$ may take negative values.

The solution for $g_H(S)$ of expression (1.2.4) for general HARA utility is

$$g_H(S) = \frac{\gamma}{\tau}A + \frac{W_0 - e^{-r_d \frac{\gamma}{\tau}}A}{\pi_H}c_H(S),$$

(e.g., Carr and Madan (1997)). The contractor wants to invest the amount $\frac{\gamma}{\tau}Ae^{-r_d}$ to the interest rate r_d and the remaining in the risky security $c_H(S)$ (if it exists). This property is usually referred to as two-fund separation. The quantity invested to the rate r_d is independent of initial wealth and the corresponding payoff $\frac{\gamma}{\tau}A$ may in some cases be interpreted as the subsistence wealth level. Since investments in the risky securities $c_H(S)$, $c_Q(S)$, and $c_L(S)$ can not yield negative payoffs, the subsistence level also represents the minimum payoff of the total investments for agents characterized by these three preferences.

If quadratic utility is assumed, the optimal payoff is

$$g_Q(S) = \frac{1}{2b} + \frac{W_0 - \frac{1}{2b}e^{-r_d}}{\pi_Q}c_Q(S). \quad (1.3.1)$$

When exponential utility is assumed we obtain

$$g_E(S) = (W_0 - \frac{1}{\tau}\pi_E)e^{r_d} + \frac{1}{\tau}c_E(S). \quad (1.3.2)$$

The investment in the risky security is decreasing with increasing absolute risk aversion τ and is independent of initial wealth W_0 .

Finally, logarithmic utility is assumed. The optimal payoff is now

$$g_L(S) = W_0c_L(S). \quad (1.3.3)$$

Observe that the market price at time zero of $c_L(S)$ is 1. In this case one-fund separation applies, i.e. the contractor wants to invest his complete initial wealth in the security $c_L(S)$.

Based on this analysis we may conclude that each of the four securities introduced earlier represents the only risky securities an investor wants to invest in (if they exist) under the four different preference scenarios.

1.3.2 Belief Specification

Both the risk-adjusted probability density $q(\cdot)$ and the contractor's belief density $f(\cdot)$ are assumed to be lognormal. A lognormal probability density with parameters m and v for the terminal stock price S is

$$\ell(m, v) = \frac{1}{\sqrt{2\pi v}S} \exp\left(-\frac{1}{2v^2} \left(\ln\left(\frac{S}{S_0}\right) - \left(m - \frac{1}{2}v^2\right)\right)^2\right).$$

We assume that $f(S) = \ell(\mu + r_f, \sigma)$ and $q(S) = \ell(r_d, \sigma)$. This choice is consistent with the set-up in continuous time by Garman and Kohlhagen (1983) which is based on the standard model⁴ by Black and Scholes (1973) and Merton (1973). Note that both densities, i.e., both the contractor and the market, agree on the volatility rate σ . Here μ can be interpreted as the contractor's belief about the expected proportional change of the currency, whereas the constant r_f is interpreted as the risk-free rate in the foreign country.

We define

$$\beta = \frac{\mu - (r_d - r_f)}{\sigma^2},$$

a constant, commonly termed Sharpe's ratio. We may interpret $\mu - (r_d - r_f)$ as the risk premium connected to investments in foreign currency, and β may be viewed as the market price of risk. We see that the expected proportional change in the exchange rate, μ , is equal to $r_d - r_f + \beta\sigma^2$, which is consistent with theories suggesting that this change is related to an interest rate differential and a risk premium, see Hodrick (1987) and Lewis (1995). Note that the assumption of a risk premium equal to zero implies that $\mu + r_f = r_d$, i.e., $q(\cdot) = f(\cdot)$.

With this complete belief specification we are able to explicitly characterize the following payoff, related to the Arrow-Debreu security and discussed earlier under equation (1.2.3),

$$\frac{f(S)}{q(S)e^{-r_d}} = CS^\beta, \quad (1.3.4)$$

where the constant $C = S_0^{-\frac{\mu - (r_d - r_f)}{\sigma^2}} \exp\left(-\frac{1}{2} \left(\frac{(\mu + r_f)^2 - r_d^2}{\sigma^2} - \mu - r_d - r_f\right)\right)$.

On the basis of this belief specification and the expressions for the optimal claims described earlier, we arrive at the following fully specified payoff functions:

$$c_Q(S) = \frac{1}{C} S^{-\beta},$$

⁴where the price process is specified as $dS_t = mS_t dt + \sigma S_t dW_t$, where dW_t represents the increment of a standard Wiener-process.

$$c_E(S) = \ln(C) + \beta \ln(S),$$

$$c_L(S) = CS^\beta.$$

By inserting these expressions into the expressions (1.3.1), (1.3.2), and (1.3.3), respectively, we derive the optimal payoff functions for our assumed beliefs. It is apparent that the numerical value of β affects the shape of the optimal payoff functions. The sign of β determines the sign of the slope, e.g., a positive β leads to optimal payoffs increasing in S . For $-1 < \beta < 0$ the optimal payoff functions are convex (and decreasing), for $0 < \beta < 1$ the optimal payoff functions are concave (and increasing). In the numerical example we have chosen a positive β less than one. According to the hypothesis of uncovered interest parity, one should not expect values of β far from zero.

Figure 1 depicts the optimal initial contingent claim for the three preference scenarios. The three utility functions are calibrated so that the coefficient of relative risk aversion (RRA) equals one⁵ for wealth level $W_0 e^{r_d}$, i.e., as all wealth was invested in the riskfree domestic asset. The three plotted g functions are similar, i.e., increasing and concave functions of roughly same shape for our choice of parameters. In addition the 45 degrees line is plotted as well as the probability density function for S under the original probability measure to indicate the likelihood of the outcomes. The numerical values of the financial parameters are assumed to be $r_d = 0.04$, $r_f = 0.06$, $\mu = 0$, and $\sigma = 0.3$, then $\beta = \frac{2}{9}$ and the constant $C = 1.0397$.

1.4 The Optimal Hedging Claim

In this section the optimal payoffs for the quadratic, exponential, and logarithmic utility functions are analyzed. Before we proceed we recall the connection between the forward price and the spot price of currency in this model given by $F_0 = S_0 e^{r_d - r_f}$.

1.4.1 Quadratic Utility

Using the quadratic utility from Table 1 and equation (1.2.7) we find by first solving for $y_Q(S)$, then determining λ from equation (1.2.6), and substituting back in the expression for $y(S)$ that

$$y_Q(S) = \frac{1}{2b} - pS + \frac{W_0 + (pF_0 - \frac{1}{2b})e^{-r_d}}{\pi_Q} c_Q(S).$$

⁵by increasing the RRA coefficient the resulting curves for the quadratic and exponential utility functions are still similar, but flatter.

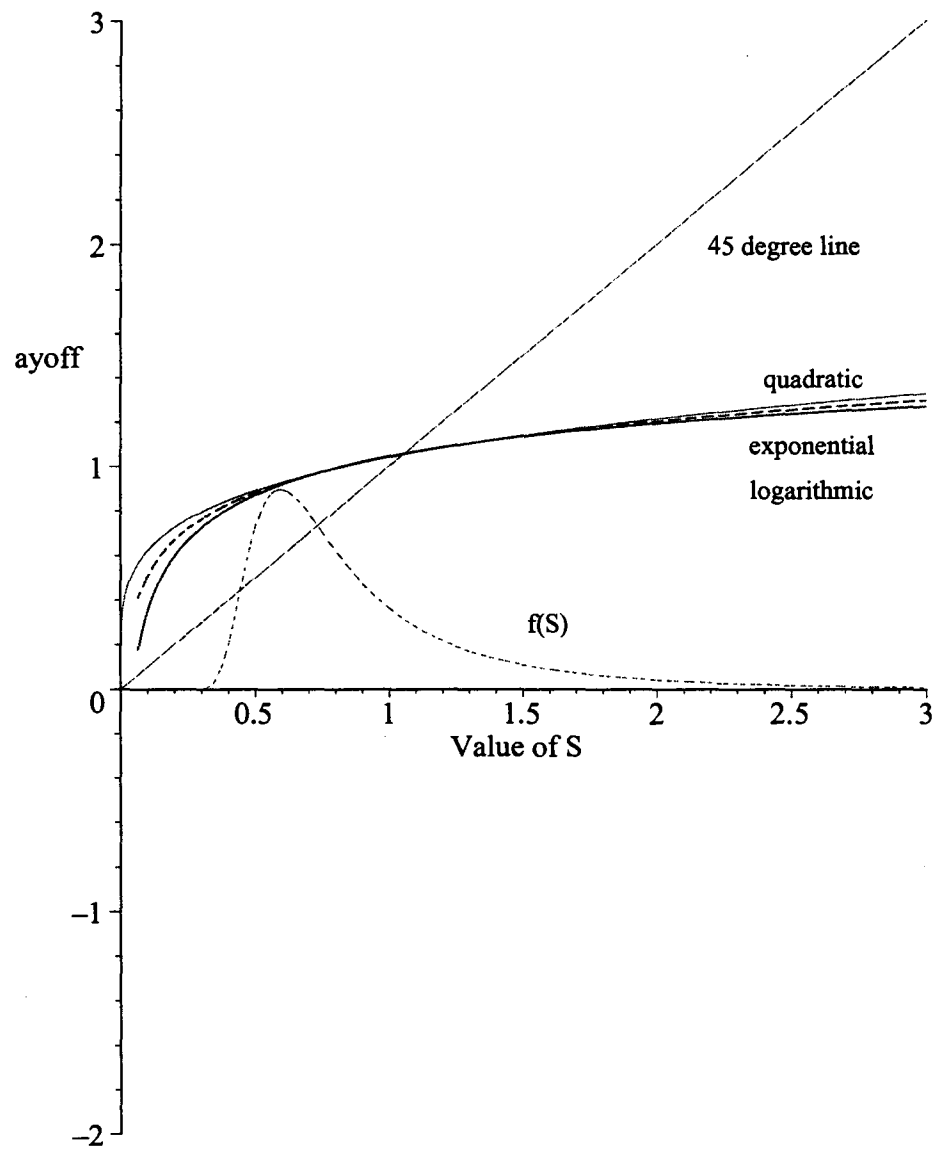


Figure 1.1: A plot of $g_Q(S)$, $g_E(S)$, and $g_L(S)$ for the parameter values $\beta = \frac{2}{9}$, $C = 1.0397$, $W_0 = 1$, $b = 0.2402$, $\tau = 0.9608$.

Two-fund separation does not longer hold since a total of three different securities are used: Riskfree investment, the security $c_Q(S)$, and the forward contract. Instead three fund separation is obtained.

The optimal hedging claim is

$$h_Q(S) = p(e^{-r_d} F_0 \frac{c_Q(S)}{\pi_Q} - S).$$

The optimal hedging strategy for quadratic utility is to sell p forward contracts on the foreign currency, borrow the present value of the forward contracts at time zero ($pF_0e^{-r_d}$) to invest in the security $c_Q(S)$.

The corresponding strategy by Steil (1993) is to sell p forward contracts. The positive risk premium introduces the asset $c_Q(S)$ which again leads to a somewhat more complicated strategy. The contractor's total investments in the security $c_Q(S)$ has increased which again influences the optimal payoff $h(S)$. In particular, $h(S)$ is a non linear function of S , which is not the case in Steil's model. Furthermore, Steil's statement that the hedging strategy does not depend on initial wealth is proved.

1.4.2 Exponential Utility

Performing the same exercise for the exponential utility function we obtain

$$y_E(S) = e^{r_d}(W_0 - \frac{\pi_p}{\tau} - \frac{\pi_E}{\tau}) + \frac{c_p(S)}{\tau} + \frac{c_E(S)}{\tau}, \quad (1.4.1)$$

where $c_p(S) = \ln(pe^{-r_d S} + 1 - p)$ and $\pi_p = \int_0^\infty c_p(s)e^{-r_d s} q(s) ds$, represents the corresponding time zero market price.

Also in this case three-fund separation applies, the first term represents the terminal value of the investment to the interest rate r_d , the last term represents the investment in the security $c_E(S)$, and the term in the middle represents the investment in the security $c_p(S)$. In contrast to the quadratic case the third fund is not the foreign currency, but a particular hedging claim with payoff $c_p(S)$.

Solving for the optimal hedging claim we obtain

$$h_E(S) = \frac{-\pi_p e^{r_d}}{\tau} + \frac{c_p(S)}{\tau}. \quad (1.4.2)$$

Observe that the higher absolute risk aversion, the lower investment in the risky securities. Also note that $h_E(S)$ only involves the hedging claim $c_p(S)$ and, in particular, is independent of the risk premium.

1.4.3 A Note on the Classic Hedging Analysis with Exponential Utility

In order to compare our results with the classic approach of confining the analysis to the optimal use of forward contracts only, we now focus on the case of a standard, non-contingent exposure, i.e., $p = 1$.

From equation (1.4.2) with $p = 1$ it follows that

$$h_E(S) = F_0 - S,$$

which is the payoff from a short forward contract. Receiving one unit of S from the project with probability one, is basically the same as increasing the initial wealth by one unit of S . Exponential utility implies that the agent has constant absolute risk aversion, so that the optimal amount invested in the risky asset should not vary with different levels of wealth. Hence, when $p = 1$, the optimal hedge, $h(S)$, should exactly offset the payoff from the project. This is obviously achieved by a forward contract.

We now denote the risk premium by $\rho = \mu - (r_d - r_f)$ so that $\beta = \frac{\rho}{\sigma^2}$, and compute for $p = 1$ from expressions (1.3.4) and (1.4.1)

$$y_E(S) = e^{r_d}W_0 + F_0 - S - \frac{\rho}{\tau\sigma^2}[e^{r_d}\pi_{\ln(S)} - \ln(S)], \quad (1.4.3)$$

where $\pi_{\ln(S)} = \int_0^\infty \ln(S)e^{-r_d}q(s)ds$.

The classic analysis differs from ours by solely focusing on the use of forward contracts. Only a brief review is given here, for more details see for instance Newbery (1989).

Within the classical approach the total payoff from a similar exposure at time 1 can be written as

$$\kappa = S + z(f - S),$$

where z is the number of forward contracts used to hedge the cashflow S .

The analysis proceeds by maximizing the certainty equivalent CE

$$CE = E[\kappa] - \frac{1}{2}\tau\sigma_\kappa^2,$$

where σ_κ^2 is the variance of κ .

The expression for CE holds exactly for exponential utility and a normally distributed exchange rate. For a lognormally distributed exchange rate and exponential utility the expression for CE can only be considered as a second order Taylor approximation.

It is straight forward to show that the optimal number of forward contracts is equal to

$$z^* = 1 - \frac{E[S] - F_0}{\tau\sigma^2}.$$

Multiplying this with the payoff from a forward contract, yields the payoff of the optimal hedge from the classical analysis, denoted $\tilde{y}(S)$:

$$\tilde{y}(S) = (F_0 - S) - \frac{E[S] - F_0}{\tau\sigma^2}(F_0 - S). \quad (1.4.4)$$

Comparing equations (1.4.3) and (1.4.4), we first note the relationship between ρ and $(E[S] - F_0)$. Recall that $F_0 = S_0 e^{r_d - r_f}$ in our model. Furthermore, $E[S] = S_0 e^{\mu}$. A risk premium ρ equal to zero, implies that $\mu = r_d - r_f$. Hence, $E[S] = F_0$ and $z^* = 1$. We see that in the case of zero risk premium, $y_E(S) = \tilde{y}(S)$.

However, $y_E(S)$ and $\tilde{y}(S)$ are different when the risk premium is positive. A positive risk premium implies that $E[S] > F_0$, and z^* becomes less than one. Hence, the optimal forward position is reduced in the classic case, but still maintaining a linear exposure to the risky asset S . Our analysis, on the other hand, shows that the optimal exposure in the hedging claim is *independent* of the risk premium. Rather, a change in the risk premium should be absorbed by a change in the payoff of the initial claim. This interesting and important result is overlooked by the restrained approach of confining the analysis to one particular hedging instrument.

1.4.4 Logarithmic Utility

Also in this case we proceed by substituting the marginal utility from Table 1 into equation (1.2.7). This leads to a quadratic expression for $y(S)$ with solution

$$y_L(S) = \frac{1}{2} \left(\frac{1}{\lambda} c_L(S) - S \right) \pm \sqrt{\frac{1}{4} \left(\frac{1}{\lambda} c_L(S) - S \right)^2 + \frac{1}{\lambda} S(1-p)c_L(S)}.$$

This can be rewritten as

$$y_L(S) = \frac{1}{2} \left(\frac{1}{\lambda} c_L(S) - S \pm \sqrt{\left(\frac{1}{\lambda} c_L(S) - (2p-1)S \right)^2 + 4p(1-p)S^2} \right).$$

From the last expression we see that it may be hard to find an analytical expression for the square-root on the right hand side for values of p different from 0 or 1.

We proceed by analyzing the case $p = 1$ and find that

$$y_L(S) = (F_0 e^{-r_d} + W_0) c_L(S) - S.$$

With help of equation (1.3.3) the optimal hedging claim is determined as

$$h_L(S) = F_0 e^{-r_d} c_L(S) - S.$$

Comparing this payoff with the corresponding one based on quadratic utility for $p = 1$, we see that both utility functions imply the same optimal strategy: Sell one forward contract, borrow the present value of the forward price at time zero, invest the borrowed amount in the risky asset, $c_L(S)$ in the logarithmic case and $c_Q(S)$ in the quadratic case.

In figure 2 the optimal hedging claims are plotted for $p = 1$. Whereas the corresponding initial claims were similar for the three utility functions, the optimal hedging claims show fundamental differences: The curve for the quadratic utility function is convex, the curve for exponential utility is linear, and the curve for logarithmic utility is concave for our choice of parameters.

The similar plot for $p = \frac{1}{2}$ is showed in figure 3. The most important difference for this case is that the curve for the exponential utility function is now convex, i.e., non-linear.




Figure 1.2: A plot of $h_Q(S)$, $h_E(S)$, and $h_L(S)$ for the parameter values $p = 1$, $\beta = \frac{2}{9}$, $C = 1.0397$, $W_0 = 1$, $b = 0.2402$, $\tau = 0.9608$.




Figure 1.3: A plot of $h_Q(S)$, $h_E(S)$, and $h_L(S)$ for the parameter values $p = \frac{1}{2}$, $\beta = \frac{2}{9}$, $C = 1.0397$, $W_0 = 1$, $b = 0.2402$, $\tau = 0.9608$.

1.5 Concluding Remarks

The prime focus of this chapter is how a non-zero risk premium affects an economic agent's optimal hedging decision. Our approach differs from that of Steil (1993), Brennan and Solanki (1981), Cox and Huang (1989), and Carr and Madan (1997) by introducing a risk premium, and by focusing on optimal contingent claims for hedging a non-marketed exposure, rather than on optimal consumption and portfolio policies.

Our main result is to describe how the size of a risk premium influence the shape of the optimal payoff, both in the case of standard exposure as well as with a contingent exposure. The payoff is characterized explicitly for three different sets of preferences. We split the optimal contingent claim into a sum of an initial component and the hedge component. In the case of exponential utility we demonstrate that the hedge component is independent of the risk premium. In case of quadratic and logarithmic utility, the hedge component exhibit similar non-linearities as the initial component. We may interpret the pure hedge both under quadratic and logarithmic utility as forward contracts plus the present value of the forward price invested in the same manner as the initial component.

Chapter 2

Optimal Multivariate Exposure: Hedging Currency Risk of Foreign Equity Investments

In this chapter we study examples of utility-maximizing, or optimal contingent claims where the payoff depends on more than one risky asset. We look at the connection between distribution of individual assets and the distribution of optimal wealth. In particular, we study optimal hedging of currency exposure from a foreign equity investment. Such hedging is motivated by the fact that currency and equity have different risk premiums. Our formulation of the problem is well suited to illustrating interaction between several risk premiums and optimal exposure. We comment on Siegel's paradox and show that it plays no part in our optimal hedge, as it does in the "universal hedging ratio" of Black (1990)

2.1 Introduction

The need to analyze and deal with simultaneous exposures to more than one risky asset or asset class arises in several situations. A standard example is an investor buying foreign equities. The total payoff as seen from the investor's perspective, is the product of the future exchange rate and the future stock price. Thus, the exposure is a function of two risky security prices.

Another way of describing a foreign equity investment is as a composite risk. It is different from holding two risky assets in a linear combination, the position is a multiplicative, or non-linear, exposure from the two assets. It resembles a situation of quantum risk: The local return on, say, the equity decides how large the currency exposure is, or vice versa. The size of your stake in one lottery depends on the outcome of another lottery. This problem is not unlike the situation of contingent exposure in chapter one, but whereas in that case there was one marketed and one non-marketed risk, we deal here with two marketed risks.

In this chapter we first find and analyze a general optimal multivariate payoff function, i.e., a function of two risky assets, that maximizes utility at a future horizon date. In other words, if an investor were not restricted to hold linear combinations of risky assets, what would the optimal payoff function from a derivative that is defined on two risky assets look like? The next question is; given that the investor are restricted to invest in a composite risk like the foreign equity investment discussed above, what utility maximizing contingent claim should be used in hedging, given that only exposure from one of the assets, e.g., equities but not currencies, is seen as favorable?

The standard mean-variance portfolio problem of Markowitz (1952) could also be seen as a multi-asset exposure problem. However, in that case a static model is assumed, and the choice variable is an optimal linear exposure in each asset, i.e., buy and hold strategies. Holdings in each asset are additively aggregated to the total optimal exposure.

As illustrated by the foreign equity example, optimal multi-asset exposure is studied in a more general sense. Our analysis is more related to the dynamic optimal portfolio choice problem formulated in Merton (1971). But rather than finding an optimal trading strategy by solving a dynamic optimization problem, we solve a static optimization problem and find an optimal payoff function at a future point in time. Cox and Huang (1989) have shown that the two approaches, i.e., finding the optimal trading strategy la Merton and finding the replicating strategy for the optimal contingent claim, yield identical solutions.

In this chapter we investigate optimal multi-asset exposure for a single individual agent. We only digress on some aggregate type issues in section 2.3.1, otherwise we do not consider equilibrium results. A similar approach can be found in Brennan and Solanki (1981), Carr and Madan (1997) and others. In this chapter, however, the model is generalized to apply to more than one risky asset.

Several of the results herein has been shown earlier, although through a different framework. In particular, a careful reading of Cox and Huang (1989) will provide many of the insights given in this chapter. However, we feel that our formulation of the pricing kernel¹ as a multivariate function of the asset prices is well suited to illustrating the interaction between the parameters of each asset's price process, e.g., the risk premium, and the optimal exposure. We think the formulation makes the analysis more accessible. Moreover, our approach provides an opportunity to discuss the link between the distribution of individual assets and the distribution

¹A *pricing kernel* is a random variable ξ such that $P = E^Q[C] = E^P[\xi C]$, where P is market value, Q a yield-equating, or risk-adjusted probability measure, C is cash-flow and P is the investor's subjective probability measure. Other popular names for this entity include *state-price deflator*, *Girsanov's factor* and *Radon-Nikodym derivative*.

of optimal aggregate wealth.

Secondly, our approach makes it quite easy to illustrate the need for derivatives of rather an exotic nature, such as quanto-derivatives. A *quanto derivative* is a derivative with a payoff that is measured in a foreign currency but pays off in the home currency. For example, a Norwegian quanto futures contract on the American S&P 500 index pays NOK 1417 when the futures price is 1417.

Thirdly, our example allows a comment on Siegel (1972)'s paradox, viz., the observation that through Jensen's inequality, the expected value of an exchange rate is not equal to the negative of the expected value of the inverse of that exchange rate. This apparently makes it difficult for two risk neutral investors, with different home currencies, to agree on an equilibrium price for the exchange rate; hence the paradox. In Black (1990), Siegel's paradox is used as an argument for why a 100% hedge of currency exposure in an equity investment is sub-optimal, regardless of the agent's belief about the expected returns on the currency or its diversifying properties. This argument is not sustained in the present analysis, because an application of the Change of Numeraire Theorem, developed in Geman, El Karoui, and Rochet (1995), essentially resolves Siegel's paradox in our model.

The chapter is organized as follows: In section 2.2 the model is introduced and a general solution is presented for HARA utility and arbitrary beliefs. In Section 2.3 we assume that two risky assets follow a two-dimensional geometric Brownian motion price process. The pricing kernel is described as a function of the two prices, and optimal contingent claims for exponential, power and logarithmic utility functions are presented. In section 2.3.1. we discuss distribution effects and section 2.3.2. comments on diversification. In section 2.4 the model is used to describe optimal hedging of currency exposure of a foreign equity investment. A comment on Siegel's paradox is given in section 2.4.1. Concluding comments are contained in section 2.5.

2.2 Optimal Exposure

In this section we develop optimal exposure as a function of asset prices at a single future date. The existence and uniqueness of an equivalent martingale measure, i.e., an equilibrium price measure or a pricing kernel, is assured by assuming no arbitrage and a dynamically complete market. The agent has increasing, concave and continuous preferences.

The model is quite similar to that of Brennan and Solanki (1981), but here we describe the

payoff as a function of several risky assets rather than as a function of the value of a reference portfolio. This allows us to study the effect of the joint distribution of the assets, not only the distribution of the aggregate, or reference portfolio, as in Brennan and Solanki (1981). To keep the analysis as simple as possible we let the agent choose between two risky assets and a risk-free alternative. Including more than two risky assets in the analysis will not provide additional insight in a general sense.

The optimal payoff function is denoted $g(\cdot)$. The whole of the investor's wealth is invested in the contract. Hence the optimal payoff is the solution to the problem:

$$\max_{g(S_1, S_2)} \mathbb{E}[U(W)] = \max_{g(S_1, S_2)} \int_0^\infty \int_0^\infty U(g(s_1, s_2)) f(s_1, s_2) ds_1 ds_2, \quad (2.2.1)$$

$$\text{subject to} \quad \int_0^\infty \int_0^\infty e^{-r} g(s_1, s_2) q(s_1, s_2) ds_1 ds_2 = W_0. \quad (2.2.2)$$

Here $W = g(S_1, S_2)$ is terminal wealth. The investor determines the optimal exposure from the market, i.e., decides the form of $g(\cdot)$, under the condition that the value of that payoff equals his initial endowment, or time zero wealth, W_0 . $U(\cdot)$ is the utility function, $f(\cdot)$ is the agent's subjective beliefs and $q(\cdot)$ is the yield-equating or risk-adjusted density function.

He decides the payoff simultaneously across all assets in the economy. Any constant in the resulting expression for the optimal $g(\cdot)$ will represent the future value of an investment in the risk-free asset.

Proceeding as in Persson and Trovik (1997) and others, we formulate the Lagrangian for (2.2.1) and (2.2.2) and take the derivative (S_1, S_2) by (S_1, S_2) , i.e., for all combinations of S_1 and S_2 to yield the first order condition

$$\lambda = U'(g(S_1, S_2)) \frac{f(S_1, S_2)}{q(S_1, S_2)e^{-r}}, \quad (2.2.3)$$

where λ is the constant Lagrange multiplier.

We now introduce a HARA utility function over final wealth:

$$U(W) = \frac{\gamma}{1-\gamma} \left(\frac{\tau}{\gamma} W - A \right)^{1-\gamma}. \quad (2.2.4)$$

By using (2.2.4), writing (2.2.3) in terms of $g(\cdot)$, inserting the resulting expression in (2.2.2), solving for λ and then using that λ in (2.2.3), we get

$$g(S_1, S_2) = \frac{\gamma}{\tau} A + \left(W_0 - e^{-r} \frac{\gamma}{\tau} A \right) \frac{c_H(S_1, S_2)}{\pi_H}, \quad (2.2.5)$$

where $c_H(S_1, S_2) = \left(\frac{f(S_1, S_2)}{q(S_1, S_2)e^{-r}} \right)^{\frac{1}{\gamma}}$ and π_H is the current price of c_H , i.e., $\pi_H = \mathbb{E}^Q[e^{-r} c_H]$. Hence the agent is allocating $\frac{\gamma}{\tau} A$ to the risk-free asset and $(W_0 - e^{-r} \frac{\gamma}{\tau} A)$ to the risky asset c_H .

Equation (2.2.5) illustrates the interaction between preferences and beliefs as a factor determining the optimal exposure. The parameters of the utility function determine how much should be invested in risky and risk-free assets respectively. The type of risky asset to choose, i.e., the form of the payoff function, is described by the pricing kernel $\frac{f(S_1, S_2)}{q(S_1, S_2)e^{-r}}$, but augmented by the $c_H(\cdot)$ function.

So far this analysis has followed the lines of Brennan and Solanki (1981), Leland (1980), Carr and Madan (1997), Persson and Trovik (1997) and others. In the following section, however, we focus on the pricing kernel in a multi-asset economy.

2.3 Optimal Multivariate Payoff when Prices are Lognormally Distributed

We adopt a generalized Black and Scholes economy with two risky securities with prices S_1 and S_2 , and a risk-free money-market account earning a constant interest rate r . We assume a dynamically complete economy. The two securities evolve according to correlated geometric Brownian motions, i.e., for $i = \{1, 2\}$,

$$dS_i = \mu_i S_i dt + \sigma_i S_i dw_i(t). \quad (2.3.1)$$

We define $\mathbf{w}(t)^\top = [\sigma_1 w_1(t), \sigma_2 w_2(t)]$, where w_i represent correlated Wiener processes. We assume the coefficients μ_i and σ_i to be constants. Furthermore, we define

$$d\mathbf{w} = \mathbf{V} dz$$

where $\mathbf{z}(t)^\top = [z_1(t), z_2(t)]$ is a vector of independent Wiener processes and where the matrix

$$\mathbf{V} \equiv \begin{bmatrix} \sigma_1 & 0 \\ \sigma_2 \rho & \sigma_2 \sqrt{1 - \rho^2} \end{bmatrix}.$$

captures the covariance between the assets. Writing out the above relation, we get

$$\begin{aligned} dw_1 &= dz_1 \\ \text{and} \quad dw_2 &= \rho dz_1 + \sqrt{1 - \rho^2} dz_2. \end{aligned} \quad (2.3.2)$$

It is easy to show that $Var(dw_i) = dt$ and $Cov(dw_1, dw_2) = \rho dt$, hence dw_i indeed represent Wiener processes.

Our aim now is to describe the pricing kernel as a multivariate function of the prices in this economy, and then use that function in the expression for the optimal payoff in (2.2.5). The pricing kernel is described in the following proposition:

Proposition 1. *In a Black and Scholes economy consisting of two risky assets, the pricing kernel as a function of the two risky asset prices can be written as*

$$\xi(S_1, S_2) = C S_1^{-\frac{A}{\sigma_1}} S_2^{-\frac{B}{\sigma_2}}, \quad (2.3.3)$$

when the horizon is normalized to one, i.e., $t = 1$. The constants A and B in the exponent is given by

$$A = \frac{1}{1 - \rho^2} \frac{\mu_1 - r}{\sigma_1} - \frac{\rho}{1 - \rho^2} \frac{\mu_2 - r}{\sigma_2}, \quad (2.3.4)$$

$$B = \frac{1}{1 - \rho^2} \frac{\mu_2 - r}{\sigma_2} - \frac{\rho}{1 - \rho^2} \frac{\mu_1 - r}{\sigma_1}, \quad (2.3.5)$$

and the constant C is given by

$$C = S_1(0)^{\frac{A}{\sigma_1}} S_2(0)^{\frac{B}{\sigma_2}} \exp \left[\left(\frac{A}{\sigma_1} \left(\mu_1 - \frac{1}{2} \sigma_1^2 \right) + \frac{B}{\sigma_2} \left(\mu_2 - \frac{1}{2} \sigma_2^2 \right) + \frac{1}{2} (\beta_1^2 + \beta_2^2) \right) \right].$$

Proof. Girsanov's factor in this economy is written as

$$\xi(t) = \exp \left[-\beta^\top \mathbf{z}(t) - \frac{1}{2} \|\beta\|^2 t \right], \quad (2.3.6)$$

where $\beta = \mathbf{V}^{-1}(\mu - r\mathbf{1})$. Hence we have

$$\beta = \begin{bmatrix} \frac{\mu_1 - r}{\sigma_1} \\ \frac{1}{\sqrt{1 - \rho^2}} \frac{\mu_2 - r}{\sigma_2} - \frac{\rho}{\sqrt{1 - \rho^2}} \frac{\mu_1 - r}{\sigma_1} \end{bmatrix} \equiv \begin{bmatrix} \beta_1 \\ \beta_2 \end{bmatrix}$$

From the solution of the stochastic differential equation (2.3.1) we have

$$w_i(t) = \frac{1}{\sigma_i} \left(\ln \frac{S_i(t)}{S_i(0)} - \left(\mu_i - \frac{1}{2} \sigma_i^2 \right) t \right), \quad (2.3.7)$$

and from (2.3.2) we have

$$\begin{aligned} z_1(t) &= w_1(t) \\ \text{and} \quad z_2(t) &= \frac{w_2(t) - \rho w_1(t)}{\sqrt{1 - \rho^2}}. \end{aligned} \quad (2.3.8)$$

Combining (2.3.6) and (2.3.8) we get

$$\xi(t) = \exp \left[- \left(\beta_1 - \frac{\beta_2 \rho}{\sqrt{1 - \rho^2}} \right) w_1(t) - \frac{\beta_2}{\sqrt{1 - \rho^2}} w_2(t) - \frac{1}{2} (\beta_1^2 + \beta_2^2) t \right].$$

To obtain Girsanov's factor written as a function of the asset prices, S_1 and S_2 , we substitute w_i from (2.3.7) in the above expression. We have assumed $t = 1$ to simplify matters. \square

We are now ready to study the combined effects of preferences as given by (2.2.4), and beliefs as given by (2.3.3). The HARA class of utility functions displays linear absolute risk tolerance. Hence, quadratic, exponential, power and logarithmic utility are special cases of HARA. Here we focus on the two viable alternatives, exponential with constant absolute risk aversion, and logarithmic and power with constant relative risk aversion.

Optimal exposure in the three cases is then found by inserting (2.3.3) into (2.2.5) and choosing parameters appropriately:

$$g(S_1, S_2)_{exp} = \left(W_0 - \frac{1}{\tau} \pi_E \right) e^r + \frac{1}{\tau} c_E(S_1, S_2) \quad (2.3.9)$$

$$g(S_1, S_2)_{pow} = W_0 \frac{c_P(S_1, S_2)}{\pi_P} \quad (2.3.10)$$

$$g(S_1, S_2)_{log} = W_0 c_L(S_1, S_2) \quad (2.3.11)$$

where

$$c_E(S_1, S_2) = \ln \left(e^r C^{-1} S_1^{\frac{A}{\sigma_1}} S_2^{\frac{B}{\sigma_2}} \right) = D + \frac{A}{\sigma_1} \ln S_1 + \frac{B}{\sigma_2} \ln S_2 \quad (2.3.12)$$

$$c_P(S_1, S_2) = \left(e^r C^{-1} S_1^{\frac{A}{\sigma_1}} S_2^{\frac{B}{\sigma_2}} \right)^{\frac{1}{\gamma}} \quad (2.3.13)$$

$$c_L(S_1, S_2) = e^r C^{-1} S_1^{\frac{A}{\sigma_1}} S_2^{\frac{B}{\sigma_2}} \quad (2.3.14)$$

and π_i for $i \in \{E, P, L\}$ is the time t market price of the claims c_i ; for example, $\pi_E = E_t^Q [e^{-r(T-t)} c_E(S_1(T), S_2(T))]$. We have $\pi_L = 1$.

Note that in this approach the object of choice is a final wealth random variable, $g(\cdot)$, rather than a trading strategy. Of course the trading strategy is implicitly chosen as the replicating strategy for $g(\cdot)$. Cox and Huang (1989) showed that this static and the more classic dynamic formulation of the optimization problem produce identical solutions.

2.3.1 Joint Distribution of Assets and Optimal Wealth

By describing the optimal payoff functions, we have implicitly described the probability distributions of optimal, terminal wealth. A study of (2.3.12), (2.3.13) and (2.3.14) together with

their corresponding $g(\cdot)$ -functions, reveals the distributions of optimal wealth.

If the agent has constant absolute risk aversion (CARA), e.g., has an exponential utility function, then the utility maximizing terminal wealth is normally distributed when prices are joint lognormally distributed. The reason is that the pricing kernel is linearized by the utility function as seen in (2.3.12).

If, however, the agent has constant proportional risk aversion (CPRA), his optimized wealth at the horizon has a lognormal distribution when prices are joint lognormally distributed. In this case the multiplicative structure in the pricing kernel is maintained also in the claims c_P and c_L as seen in (2.3.13) and (2.3.14).

In this paper, we are analyzing an arbitrary individual, he is not necessarily representative, and he is assumed to be infinitesimal relative to the rest of the economy. The analysis performed is partial. The observations above, however, provide an opportunity to digress to issues concerning equilibrium price processes. If we for the rest of this section think of the utility functions as representative agents, and the optimal claims describing aggregate payoffs, then the observations above provide a link between assumptions concerning the distribution of single assets, assumed preferences and the distribution of optimal aggregate wealth.

Our last observation, regarding (2.3.13) and (2.3.14), aligns nicely with Bick (1987) and Bick (1990), where it is found that if the market portfolio, i.e., aggregate wealth, follows a geometric Brownian motion, then necessarily the utility function of the "representative individual" exhibits constant proportional risk aversion. Furthermore, in Stapleton and Subrahmanyam (1984) it is shown that a necessary and sufficient condition for a risk-neutral valuation relationship (RNVR) to exist, given an assumption of a joint, multivariate log-normal distribution for the underlying prices and aggregate wealth, is CPRA utility.

It is more difficult to align our observation regarding CARA utility with the corresponding issue in Stapleton and Subrahmanyam (1984). They show that under the assumption of an arbitrary, multivariate normal distribution of the underlying prices and aggregate wealth, a necessary and sufficient condition for a RNVR to hold, is that the representative investor exhibits CARA. Our observation from (2.3.12) might indicate that a joint normal distribution of assets and aggregate wealth is not a proper assumption, given a representative investor with CARA utility.

Rather, (2.3.12) provides motivation for assuming a semi-lognormal distribution, i.e., that aggregate wealth and each individual asset is bivariate normal-lognormal distributed. Such a distribution is described by Crow and Shimizu (1988), and introduced to finance by Camara

and Stapleton (1998). They study sufficient conditions for a RNVR to hold under CPRA or CARA, assuming a joint normal-lognormal distribution for aggregate wealth and asset prices in both cases. Their motivation for introducing the normal-lognormal distribution in their paper, is made with reference to the fact that the lognormal distribution is not stable under addition. Hence, if asset prices are lognormal distributed, then aggregate wealth, seen as a linear combination of assets, cannot be lognormal.

Our observations, on the other hand, indicate that the distribution of aggregate wealth varies with the assumed preferences of the representative agent, for the same assumption regarding the distribution of individual asset prices. Hence, rather than exogenously assume both a joint distribution of prices and aggregate wealth, as well as preferences of a representative agent, the distribution of aggregate wealth should follow endogenously from assumptions on the distribution of prices and preferences of a representative agent.

It should be noted that we use the term "aggregate wealth" as describing the optimal contingent claims given above. The term is used similarly in Stapleton and Subrahmanyam (1984), where the stochastic part of aggregate wealth is the end-of-period payoffs associated with contingent claims defined on the underlying assets (see Stapleton and Subrahmanyam (1984), p. 210).

Regarding the stability under addition issue: We see from e.g (2.3.14) that both asset prices and aggregate wealth are lognormal, when assuming CPRA. This apparent paradox disappears when realizing that the contingent claim payoff, or exposure, in (2.3.14) is a result of a dynamically traded portfolio consisting of linear combinations of the underlying assets. The effect of the trading (i.e., the replication of the contingent claim) makes aggregated wealth lognormal as well. This is in line with standard dynamic portfolio results from for example Merton (1971). There, the "two fund separation theorem" is described in a dynamic setting, and with lognormal asset prices. Merton (1971) shows, in a model with lognormal distributed prices, that the risky fund (the portfolio) has a lognormal distribution as well, because of similar effects.

2.3.2 Diversification

The formulae in (2.3.9) to (2.3.14) also provide some insight into the demand for risky assets. We know from standard analysis (see Huang and Litzenberger (1988)), that if an agent is long (short) a risky asset, then there exists at least one asset with a positive (negative) risk premium. If no risky assets are held, then the risk premium is zero for all risky assets.

If the risk premiums in (2.3.4) and (2.3.5) are altered, the same result is obtained in this model. If both μ_1 and μ_2 are equal to the risk free rate, then $A = B = 0$. In that case the optimized $g(\cdot)$ does not depend on S_1 nor S_2 . Hence, the entire wealth is invested in the risk free asset. If either μ_1 or μ_2 are different from r , then $g(\cdot)$ will depend on at least one of the risky assets.

Note that when one of the risk premiums is zero *and* the risky assets are uncorrelated, we see from (2.3.4) or (2.3.5) that only the risky asset with a non-zero risk premium is held in optimum. Thus, no cross-sectional diversification occurs in this case. However, if the correlation is non-zero, then both assets will be held even if one of the risk premiums is zero.

A similar result is found in the static mean-variance framework. In a mean-variance framework, the efficient frontier of all assets is a line joining the risk free rate and the tangency portfolio. The tangency portfolio is a portfolio on the efficient frontier of risky assets that maximizes the slope of that line.

In a two-asset example, suppose one of the assets has a zero risk premium. If the correlation between the assets is less than zero, the tangency portfolio contains both assets in proportions of less than 100%. If the correlation is zero, the tangency portfolio contains only the asset with the highest expected return. When the correlation is positive, again both assets are included in the tangency portfolio, but now the asset with the highest expected return is leveraged by shorting the asset with zero risk premium. Hence, the weight of the high return asset is greater than 100%. There exists a certain threshold level for the correlation, depending on expected returns, the covariance matrix and the level of the risk free rate, above which the tangency portfolio is held short, i.e., the signs of the portfolio weights just described are reversed.

Our result differs from the standard (static) mean-variance approach in that the payoff function is not linear in either the risky asset or in the tangency portfolio. This of course is due to our more general model where non-linear payoffs are allowed, or in other words where a dynamic strategy can be followed, replicating the $g(\cdot)$ payoff.

2.4 Hedging Currency Risk of Foreign Equity Exposure

In this section we use the model laid out above to study an optimal hedging problem. Suppose an investor invests in equities globally. He invests in physical stocks, and thus gains a currency exposure implicit in the equity investment. We assume his views, i.e., return expectations, may differ for the two asset classes. This means that the exposure in the two asset classes should

differ as well, so that a need for hedging (or speculation) is implied. A hedge may decouple the two asset exposures so that they can be evaluated separately. We study how such a hedge can be structured optimally in a utility framework.

To ensure consistency in an optimal utility setting, one cannot assume that the framework is only used partially on the investor's wealth. Hence, the initial, linear investment in the foreign stock we assume is made for some e.g. institutional reason outside our model.

We interpret S_1 as representing the price of the stock in foreign currency, and S_2 as the price of the foreign currency in home currency terms. If the investor was permitted to maximize his utility straight away, the optimal exposure in foreign equity and currency is already described in a general sense by $g(\cdot)$ in (2.2.5). However, a presence of a linear investment in the foreign stock alters the problem somewhat. The problem can be formulated in the following way:

$$\max_{y(S_1, S_2)} \mathbb{E}^{\mathbb{P}} [U(y(S_1, S_2) + S_1 S_2)] \quad (2.4.1)$$

$$\text{s.t.} \quad \mathbb{E}^{\mathbb{Q}} [e^{-r_d} (y(S_1, S_2) + S_1 S_2)] = W_0. \quad (2.4.2)$$

Here, the investor first invests linearly, i.e., buys and holds, in the foreign stock. Then he invests the remaining initial wealth in a claim paying off $y(\cdot)$, taking into account the payoff from his linear investment. The optimal payoff in the problem above, $y(S_1, S_2)$, may differ from $g(S_1, S_2)$ found in (2.2.5) due to the presence of the linear investment. We define this difference, i.e., $y(\cdot) - g(\cdot)$, as the optimal hedge; $h(\cdot)$. In other words, the hedging contract, $h(\cdot)$, together with the linear investment in the foreign equity that is being hedged, produce the exposure the investor otherwise would have had.

To adapt the model to the case of one foreign equity and one exchange rate, i.e., two different numeraire currencies or money market instruments, some modifications of the β in (2.3.6) is necessary. For the currency, we modify β_2 by letting $r = r_d$, i.e., the domestic interest rate, and include the payout rate for the currency by letting $\mu_2 = \hat{\mu}_2 + r_f$, where $\hat{\mu}_2$ is the expected appreciation rate. For the stock, we modify β_1 by letting $r = r_f$, i.e., the foreign interest rate, and by letting $\mu_1 = \hat{\mu}_1 + \rho\sigma_1\sigma_2$. We assume no dividends from the stock. With these parameters in (2.3.6), an application of Girsanov's theorem produces the following transformations of the processes (2.3.1) under the risk-adjusted measure for the home numeraire:

$$dS_1 = S_1 \left[(r_f - \rho\sigma_1\sigma_2)dt + \sigma_1 dw_1^{\mathbb{Q}} \right] \quad (2.4.3)$$

$$dS_2 = S_2 \left[(r_d - r_f)dt + \sigma_2 dw_2^{\mathbb{Q}} \right] \quad (2.4.4)$$

We know that in a complete economy, absence of arbitrage is equivalent to existence of a unique martingale measure for a given numeraire, i.e., there is one measure for each numeraire such that all discounted securities rebased in that numeraire are martingales. Under the measure \mathbb{Q} , the discounted gains process for S_2 is a martingale, but S_1 is not. The reason for that is that S_1 is denominated in the foreign numeraire, and \mathbb{Q} is only a martingale measure for the home numeraire. When the price of the foreign equity is rebased to the home numeraire by multiplying with the exchange rate, S_2 , we see after an application of Ito's lemma that the discounted process for $S_1 S_2$ is a martingale under \mathbb{Q} :

$$\begin{aligned} d(S_1 S_2) &= S_1 S_2 \left[(r_f - \rho \sigma_1 \sigma_2 + r_d - r_f + \rho \sigma_1 \sigma_2) dt + \sigma dw^{\mathbb{Q}} \right] \\ &= S_1 S_2 \left[r_d dt + \sigma dw^{\mathbb{Q}} \right], \end{aligned} \quad (2.4.5)$$

where $\sigma = (\sigma_1^2 + 2\rho\sigma_1\sigma_2 + \sigma_2^2)^{\frac{1}{2}}$.

The other risk-adjusted probability measure in the present case, denoted \mathbb{Q}_f , applies to payoffs in the foreign numeraire and may be interpreted to represent risk neutral expectations seen from the foreign investor's perspective. Under this measure the discounted gains processes for S_1 and $\frac{1}{S_2}$ must be martingales to rule out arbitrage. Also, we must have that for any discounted payoff in either numeraire, X^d or X^f ,

$$E^{\mathbb{Q}}[X^d] = S_2 E^{\mathbb{Q}_f}[X^f],$$

i.e., the value calculated must of course be independent of the choice of numeraire.

By reverse engineering the Girsanov theorem, we see from (2.4.3) that

$$dw_1^{\mathbb{Q}_f} = dw_1^{\mathbb{Q}} - (\text{cov}(dS_1, dS_2)/\sigma_1) dt \quad (2.4.6)$$

Using (2.4.4) and Ito's lemma on $d(\frac{1}{S_2})$ we have that

$$d\frac{1}{S_2} = \frac{1}{S_2} \left[(r_f - r_d + \sigma_2^2) dt + \sigma_2 dw_2^{\mathbb{Q}} \right]. \quad (2.4.7)$$

Again we can see what the $\frac{d\mathbb{Q}}{d\mathbb{Q}_f}$ must look like through Girsanov's theorem:

$$dw_2^{\mathbb{Q}_f} = dw_2^{\mathbb{Q}} - (\text{cov}(d(1/S_2), dS_2)/\sigma_2) dt. \quad (2.4.8)$$

In our lognormal model we have that $\text{cov}(d(1/S_2), dS_2) = \text{cov}(-\ln S_2, \ln S_2) = -\sigma_2^2$. The covariance terms above are in the market known as "convexity corrections" (see p. 142 Pelsser (2000)), and are applications of the Change of Numeraire Theorem which was introduced by Geman, El Karoui, and Rochet (1995). The desire to evaluate quanto derivatives

gave attention to the effect a change of numeraire has on a price process. Reiner (1992) showed that (2.4.3) is the correct specification when pricing a derivative which is measured in the foreign equity price put pays off in home currency, i.e., a quanto derivative. For other numeraire considerations see e.g. Hull (2000), Pelsser (2000) and Musiela and Rutkowski (1997).

We can re-formulate the problem in (2.4.1) and (2.4.2) so that we solve for the hedging claim only. From (2.4.5) we see that

$$\mathbb{E}_t^{\mathbb{Q}} \left[e^{-r_d(T-t)} S_1(T) S_2(T) \right] = S_1(t) S_2(t),$$

and the condition (2.4.2) can be written as

$$\mathbb{E}^{\mathbb{Q}} \left[e^{-r_d} y(S_1, S_2) \right] = \hat{W}_0, \quad (2.4.9)$$

where $\hat{W}_0 = W_0 - S_1(t)S_2(t)$, i.e., the remaining initial wealth after the linear investment in foreign equity has been made. Because \hat{W}_0 is optimally invested in $g(\cdot)$, i.e., $\mathbb{E}^{\mathbb{Q}} [e^{-r_d} y(S_1, S_2)] = \mathbb{E}^{\mathbb{Q}} [e^{-r_d} (g(S_1, S_2) + h(S_1, S_2))] = \hat{W}_0 + \mathbb{E}^{\mathbb{Q}} [e^{-r_d} h(S_1, S_2)]$, the problem in (2.4.1) and (2.4.2) can be equivalently formulated as

$$\max_{h(S_1, S_2)} \mathbb{E}^{\mathbb{P}} [U(h(S_1, S_2) + S_1 S_2)] \quad (2.4.10)$$

$$\text{s.t.} \quad \mathbb{E}^{\mathbb{Q}} [e^{-r_d} h(S_1, S_2)] = 0. \quad (2.4.11)$$

Note that this is a problem in one function, $h(\cdot)$. Hence, it is total exposure that is hedged, the exposure in currency and equity is not treated as separate decision variables.

The solution is readily available through the same procedure as in (2.2.1) to (2.2.5). To exemplify a solution, we now assume a power utility function. Then the optimal payoff from the hedge is

$$\begin{aligned} h(S_1(T), S_2(T)) &= \frac{S_1(t)S_2(t)}{\pi_P} (e^{-r_d} \xi(S_1, S_2))^{-\frac{1}{\gamma}} - S_1(T)S_2(T) \\ &= \frac{S_1(t)S_2(t)}{\pi_P} c_P - S_1(T)S_2(T), \end{aligned} \quad (2.4.12)$$

where c_P is given in (2.3.13). This payoff can be replicated by means of the following transactions: Sell $S_1(T)S_2(T)$ forward and borrow the present value of the forward price, $S_1(t)S_2(t)$. That amount is enough to invest in $\frac{S_1(t)S_2(t)}{\pi_P}$ units of the contingent claim c_P . Repay the loan with the forward price received at T . As indicated earlier, the hedge is a contingent claim that changes the exposure from the initial linear exposure to the optimal exposure as given in (2.3.13).

The amount of currency exposure (or hedging) can be seen from the extent c_P depends on the currency S_2 . This, as discussed earlier, depends on the correlation between S_2 and S_1 as well as on the expected appreciation rate (change in exchange rate). If $\hat{\mu}_2 = r_d - r_f$, i.e., the expected appreciation is equal to the interest differential, then currency exposure is held only if the correlation between the currency and the stock is different from zero. This can be seen from (2.3.13) and the expressions (2.3.4) and (2.3.5). The intuition behind this is very much in line with the standard mean-variance argument, as pointed out earlier.

Note however, that $h(\cdot)$ is defined as a payoff in home currency units. That does not change if S_2 drops out of the equation because of a belief in a zero risk premium and a zero correlation. The optimal exposure in the foreign equity is thus a claim that pays off in home currency units. Hence, the optimal investment in this case is in a quanto derivative defined on foreign stock.

2.4.1 A Comment on Siegel's Paradox

Siegel (1972)'s paradox is an example of Jensen's inequality applied to the inverse of a stochastic variable. It is a fact that the expected value of a variable, e.g., a price of a currency, is not equal to the negative of the expected value of the inverse of that variable, seemingly violating an equilibrium condition for two risk neutral investors entering into opposite sides of a currency trade, i.e., short/long currency positions.

In our model, we see this effect by comparing the trend part in (2.4.4) with (2.4.7). If μ_{ij} is investor i 's mean (expected) return from currency j and μ_{ji} is investor j 's mean (expected) return from currency i , we have that $\mu_{ij} + \mu_{ji} = \sigma_2^2$ prior to the convexity correction, or change of numeraire, in (2.4.8).

The convexity correction in (2.4.8) (and in (2.4.6) for that matter) is a result of a no arbitrage requirement, as shown in the Change of Numeraire Theorem of Geman, El Karoui, and Rochet (1995). It has the effect of making $\mu_{ij} = -\mu_{ji}$ in a risk-neutral world, i.e., under the measures \mathbb{Q} and \mathbb{Q}_f .

Several attempts have been made at explaining Siegel's paradox. Accepting it implies a belief that there is no equilibrium in currency markets, as indicated in Dumas, Jennergren, and Næslund (1995) p.459:

It has long been understood that the Siegel (1972) paradox is not truly a paradox but simply a result of the non-existence of equilibrium. When two risk-neutral investors use different price deflators to make investment decisions, no equilibrium

exist in the financial market. Such is the case in an international capital market when Purchasing Power Parity (PPP) does not prevail.

The role of PPP in this setting was first seen by Boyer (1972) as indicated by McCulloch (1975). When PPP holds, profits (in expectations) cannot be made from Siegel's paradox.

Even though Reiner (1992) used the same reasoning as that in the Change of Numeraire Theorem in his pricing of quanto-derivatives, the theorem was first presented in Geman, El Karoui, and Rochet (1995). In Hull (2000) the theorem is used to explain Siegel's paradox. Hints to the same line of arguing, i.e., acknowledging that there is two expectations, can also be found as early as Roper (1975) and to certain degree in Siegel (1975). In the latter reference Siegel plays down the importance of his finding:

... if medians (rather than expectations) are used ... , then the whole "paradox" completely disappears. However, as McCulloch (1975) has shown, all these problems should concern only the pure theorist, since the distortions caused by applying Jensen's Inequality to an arbitrary choice of numeraire are so small as to be empirically insignificant.

That view is contradicted in Black (1989) and Black (1990). In Black (1990) it is found that Siegel's paradox makes investors want to hold a positive amount of exchange risk. He assumes a world similar to that of the capital asset pricing model, e.g., normally distributed asset prices. When risk tolerances are the same across currencies, he finds that each investor currency hedges each investment by the fraction $(\mu_m - \sigma_m^2)/(\mu_m - \frac{1}{2}\sigma_e^2)$, where μ_m is the average, across investors, world market portfolio expected excess return, σ_m^2 is the average world market portfolio excess return variance and σ_e^2 is the average exchange rate variance. The fraction hedged does not depend on any specific means, variances or covariances except through the averages above. The sole driver behind this result is Siegel's paradox, which is behind the equation $\mu_{ij} + \mu_{ji} = \sigma_{ij}^2$. Black (1990) shows that under the assumption $\mu_{ij} + \mu_{ji} = 0$, i.e., if Siegel's paradox disappeared, the universal hedging ratio becomes 100%. In that case currency should only be held in accordance with the investor's view on the expected return, or its diversifying properties.

Black (1989) is a more practically oriented paper where the "universal hedging ratio" is estimated to 77%. That hedging ratio is thought of as a base case to be applied if the investor have no special views (i.e., adhere to consensus) on the currency (p.22, Black (1989)):

When you have special views on the prospects for a certain currency, or when

a currency's forward market is illiquid, you can adjust the hedging positions that the formula suggest.

The "universal hedging ratio" is an integral part of the Black-Litterman model, an asset allocation model widely used throughout the investment community, and advocated by a large investment bank. Thus, the empirical significance of Siegel's paradox can be substantial, contrary to Siegel's belief.

The hedging-result obtained in this note does not contain any trace of a Siegel effect. The only reason to hold any currency exposure is either based on the investor's expectations, i.e., on $\hat{\mu}_2$, or because of a diversification effect when ρ is different from zero. This can clearly be seen from (2.3.5). Siegel's paradox disappears as a result of the no arbitrage assumption and the application of the Change of Numeraire Theorem.

2.5 Concluding Remarks

In this chapter we have developed an expression for the pricing kernel as a function of prices, in a two-asset economy. We have used the expression to illustrate optimal multi-variate investing, as well as a means of analyzing optimal hedging of the currency exposure implicit in a foreign investment.

Siegel's paradox has had a long life, even after quanto assets has been thought of and priced correctly. Several papers like Briys and Bruno (1992), Dumas, Jennergren, and Næslund (1995) and Bardhan (1995), contain frameworks that display Siegel's paradox. This illustrates the importance of the Geman, El Karoui, and Rochet (1995) paper, and the Change of Numeraire Theorem.

In this paper we show that the currency exposure of an equity investment should be hedged 100% if you do not have a specific view on the currency, i.e., you adhere to the view that the forward rate is your best guess on the future exchange rate, and if the correlation between the equity and the currency is zero. If the correlation is not zero, you should bare some currency risk.

A 100% hedge in our model means that the hedging instrument is a derivative that removes, or immunize, the joint equity and currency exposure from the initial linear investment in foreign stock, and gains a pure equity exposure with payoff in the home numeraire as in a quanto derivative defined on foreign stock.

Chapter 3

Efficient Replication of Derivatives Defined on Portfolios

In this chapter we propose a method to efficiently replicate a perpetual put option on a buy-and-hold portfolio. Enhanced efficiency is achieved by letting the delta for each constituent asset of the portfolio be dependent on that single asset's volatility. Standard approximations often compute the delta on the basis of the volatility of the portfolio. The proposed method is tested with Monte Carlo simulations and compared with a standard approach based on a geometrically weighted portfolio. The proposed strategy compares favorably by displaying a decreased turnover and yielding a similar or slightly better performance. The method is particularly suitable when the portfolio assets have large differences in volatility, for instance in asset allocation settings.

3.1 Introduction

Given the current popularity of a risk-measure such as value-at-risk, it is fair to assume that downside risk is a widespread concern in the capital management industry. To avoid downside risk, there are roughly two different alternatives open to the investor. Either he can follow a stop-loss policy, i.e., cancel the exposure after a loss of a pre-specified magnitude has occurred, or he can purchase a hedge such as a put-option on his exposure to avoid such a loss. In this chapter we propose and analyze a dynamic hedging strategy that aims at replicating an option at low cost.

In the case where the exposure is a portfolio, traded options on that particular portfolio are not likely to be found. Probably a dynamic hedging strategy, replicating the option, is required. Moreover, most investors have an indefinite horizon for their portfolio. Most derivatives, however, have a fixed maturity date. To reconcile this discrepancy, costly rolling

of the contracts is necessary.

In this chapter we propose and test a method to dynamically replicate a floor on a buy-and-hold portfolio. We focus on replicating options with no maturity, i.e., time to maturity is infinite. We limit our focus in this chapter to portfolios of only long or only short positions. We assume that the portfolio is not easily represented by a traded index.

Despite their widespread popularity, buy-and-hold portfolios are difficult to reconcile with utility-based models. In this chapter, however, we concentrate on developing efficient methods for implementation, rather than discussing what investment policy should be followed. Optimal portfolio insurance has been studied in an utility-based setting by Grossman and Vila (1989) among others. They modify the Merton (1971) problem by including a minimal value for the portfolio as an additional constraint in the optimization problem.

A common result in this line of models is that a constant optimal fraction of total wealth is to be held in one or more risky assets. To follow such a rule, frequent trading is of course necessary. If a portfolio is continuously re-balanced in order to keep relative weights constant, assets that increase in value relatively more than the portfolio itself will be sold, while assets that increase less than average are bought. Such a re-balancing scheme would constitute a trading strategy resulting in a concave payoff schedule for assets with relatively larger price movements, i.e., greater volatility, than that of the average asset. The conditional variance of such a portfolio is constant through time, and from Merton (1971) we know that its value is lognormally distributed if each single asset has a lognormal distribution.

On the other hand, the value of a portfolio with constant holdings, i.e., a buy-and-hold portfolio, has very different distributional properties. Assuming a lognormal distribution for asset prices, a portfolio with constant absolute weights has unknown distribution because the sum of lognormally distributed random variables has unknown distribution, and is specifically *not* lognormally distributed. Moreover, the conditional variance of the portfolio is stochastic even if the volatility of each individual asset is constant. This follows from the fact that the relative weights change when the relative prices change and the number of assets is kept constant. Hence, the volatility of each asset contributes to the volatility of the portfolio with changing weights.

Because of the unknown distribution of the portfolio, standard methods of pricing and analyzing options on such a portfolio are not applicable.

A common approximation is simply to treat the portfolio *as if* it was lognormally distributed, and calculate the holdings in the replicating strategy from a standard option-pricing

formula. This approach leads to the holdings of each asset being calculated on the basis of the volatility of the portfolio. Thus, replicating holdings of each asset only differ because of their different relative share in the portfolio, and not because of differences in their stochastic properties.

The method we propose, on the other hand, takes into account the volatility of each asset when the hedge-ratio for that asset is calculated. The method is quite general and does not depend on the particular characteristics of the option.

We chose here to test the method on an American put option with an infinite horizon, often called a perpetual put option. In addition to its favorable characteristics in relation to risk management, as outlined above, this is a very simple option to analyze analytically. Its closed form analytics were described as early as in Samuelson (1965).

The chapter is organized as follows. Section 3.2 gives an outline of a standard pricing model and provides some well known results for later reference. Section 3.3 describes the problem and motivates an approximated solution. In section 3.3.1 we reproduce a very quick and dirty approach to the approximation. In section 3.3.2 we propose a new approximation and in section 3.3.3 the geometric approach is outlined. Our proposed strategy is compared with the geometric approach using Monte Carlo simulations in section 3.4. Section 3.5 contains concluding remarks.

3.2 The Financial Market

We consider a standard financial market model consisting of d risky assets and one risk-free bank account with constant interest rate r . The market is open for trade in the time interval $[0, \infty)$. We assume that the prices of risky assets evolve according to possibly correlated geometric Brownian motions, denoted GBM. Hence, on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ the $d \times 1$ price vector $\mathbf{S}(t)$ of risky assets evolves according to

$$d\mathbf{S}(t) = \mathbf{I}_S (\mathbf{m}dt + \mathbf{V}d\mathbf{w}(t)).$$

Here, \mathbf{I}_S is the $d \times d$ diagonal matrix of $S_i(t)$'s, \mathbf{m} is the $d \times 1$ vector of expected returns with elements μ_i , \mathbf{V} is a $d \times d$ matrix such that $\mathbf{V} \cdot \mathbf{V}^T = \mathbf{C}$ is the variance-covariance matrix with elements σ_{ij} , and finally, $\mathbf{w}(t)$ is a $d \times 1$ vector of independent Wiener processes. We will take μ_i and σ_{ij} to be constants, thereby assuming a stationary process and lognormal distribution for all risky assets.

We know from standard theory of finance that in a dynamically complete economy with no arbitrage, there exists a unique equivalent martingale measure Q . Under this measure all expected returns are equal to the risk-free rate r ; hence there exists a standard Brownian Motion \mathbf{w}^Q in R^d such that

$$dS(t) = \mathbf{I}_S (r\mathbf{1}dt + \mathbf{V}d\mathbf{w}^Q(t)).$$

From Girsanov's theorem we know that $d\mathbf{w}^Q = d\mathbf{w} - \beta dt$, where $\beta = -\mathbf{V}^{-1}(\mathbf{m} - r\mathbf{1})$. Furthermore, the Radon-Nikodym derivative

$$\xi_t = \exp\left(\int_0^t \beta(s)^T d\mathbf{w}(s) - \frac{1}{2} \int_0^t |\beta(s)|^2 ds\right)$$

simplifies to $\xi_t = \exp(\beta^T \mathbf{w}(t) - \frac{1}{2} |\beta|^2 t)$ when μ_i and σ_{ij} are assumed to be constant. For any random variable Z we have

$$\mathbb{E}_t^Q[Z] = \frac{1}{\xi_t} \mathbb{E}_t^P[\xi_T Z]$$

In a complete financial market, the value of any uncertain payoff Z can be calculated as an expectation under a risk-adjusted probability measure. These well known characteristics of the model are recorded here for later reference. Further details of the model can be found in Duffie (1996).

3.3 The Problem

In the stated model, the value of a portfolio is the sum of lognormally distributed asset prices. We are studying a case where the portfolio manager is following a buy and hold strategy; hence the holdings in the portfolio are constants. We denote the value of the portfolio at time t by $W(t)$ such that

$$W(t) = \sum_{i=0}^d n_i S_i(t), \tag{3.3.1}$$

where n_i is the number of units of asset i held in the portfolio.

The task at hand is to design a trading strategy, involving the assets in the portfolio and a risk-free asset, to provide a floor for the value of the portfolio. We want the strategy to be time invariant because the investor's aversion to outcomes below the floor is not related to any particular point in time. The perpetual American put option provides the desired

payoff structure. In the following some standard formulae are presented, serving as a point of departure for developing the proposed trading strategy.

The value of a put option of American type may be written as

$$p^*(W(t)) \equiv \sup_{\tau \in \mathcal{T}(t)} \mathbb{E}_t^{\mathbb{Q}} [e^{-r\tau} (K - W(\tau))^+] = \sup_{\tau \in \mathcal{T}(t)} \mathbb{E}_t^{\mathbb{Q}} \left[e^{-r\tau} (K - \sum_{i=1}^d n_i S_i(\tau))^+ \right] \quad (3.3.2)$$

where $\mathcal{T}(t)$ denotes the set of stopping times valued in $[t, \infty)$. This set is what defines this option as an infinitely lived one; a standard American option would have a finite endpoint.

The optimal stopping time, τ^* , is defined as the first exit time from the *optimal continuation region*, denoted D :

$$\tau^* = \inf\{t > 0 : \mathbf{S} \notin D\}.$$

The optimal continuation region is defined as the set in time-state space such that the reward from killing the process is now less than the expected reward from killing at the optimal stopping time. When the reward function is time separable and the time horizon is infinite, the continuation region is independent of t . Hence we define

$$D = \{\mathbf{S} : p(W(t)) \leq p^*(W(t))\},$$

where $p^*(W(t))$ is defined as above and $p(W(t))$ is the reward from exercising the option immediately, i.e.,

$$p(W(t)) = (K - \sum_{i=1}^d n_i S_i(t))^+.$$

Conversely, we define the *optimal stopping region* as

$$\neg D = \{\mathbf{S} : p(W(t)) \geq p^*(W(t))\}.$$

Hence we have that

$$p^*(W(t)) = \sup_{\tau \in \mathcal{T}(t)} \mathbb{E}_t^{\mathbb{Q}} [e^{-r\tau} p(W(\tau))] = \mathbb{E}_t^{\mathbb{Q}} [e^{-r\tau^*} p(W(\tau^*))].$$

For further details of optimal stopping problems see Øksendal (1998).

We will refer to $p^*(W(t))$ as the *original option* and to the corresponding stopping problem as the *original problem*. This problem is easily solved when $W(t)$ is lognormal as in Samuelson (1965). However, in the present case $W(t)$ has unknown distribution because it is the sum of lognormals. Hence, the solution must be approximated in some way. A common solution is to simply treat $W(t)$ as a single lognormal variable. In such an approach the replicating

strategy is in effect computed from one risky asset with “average” stochastic properties of $W(T)$. The holdings in each single asset are then computed from the partial derivatives of $W(T)$. However, in this approach degrees of freedom are lost. Any dependencies between the stochastic properties of $W(T)$ and *how* that value is made up of the individual values of the assets are ruled out by the model.

In a later section we investigate an approach of this kind to compare with the model of the following section. In the following we first reproduce one standard way of approximating the value of the option and see how the hedge ratios are computed from that approach. Then we aim at incorporating the characteristics of the single assets in the computation of the portfolio deltas.

3.3.1 A Standard Approximation

Samuelson (1965) showed that the value and the optimal exercise policy of a perpetual put are time-invariant. Thus, because the optimal exercise price is a constant and not a function of time, as in a standard American option, this American option value has a closed form expression.

A perpetual put with strike K on a single, lognormally distributed asset $S(t)$, with volatility σ , has the following value:

$$p^*(S) = \begin{cases} \left(\frac{K}{\gamma+1}\right) \left(\frac{(\gamma+1)S(t)}{\gamma K}\right)^{-\gamma} & \text{if } S(t) > KC, \\ K - S(t) & \text{if } S(t) \leq KC, \end{cases} \quad (3.3.3)$$

where r is the risk-free interest rate, and the exercise price is defined by the strike and a constant $C = \frac{2r}{2r+\sigma^2}$, and $\gamma = \frac{2r}{\sigma^2}$. The hedge ratio, or “delta”, for this option is found by computing $\frac{\partial p^*(S)}{\partial S}$. We have

$$\frac{\partial p^*(S)}{\partial S} = \begin{cases} -\frac{\gamma}{S} \left(\frac{K}{\gamma+1}\right) \left(\frac{(\gamma+1)S(t)}{\gamma K}\right)^{-\gamma} & \text{if } S(t) > KC, \\ -1 & \text{if } S(t) \leq KC. \end{cases}$$

A standard way of approximating the value of an option on a portfolio is to falsely assume that the portfolio is lognormally distributed. If that were the case, the value of the option could simply be calculated by replacing $S(t)$ with $W(t)$ in (3.3.3) above. The volatility is the volatility of the portfolio calculated in a standard way.

To find the delta of the option with respect to the i th asset in the portfolio, we compute $\frac{\partial p^*(W)}{\partial W} \frac{\partial W}{\partial S_i}$. From (3.3.1) we see that the derivative of W with respect to the value invested in

each asset is equal for all assets, i.e., $\frac{\partial W}{\partial n_i S_i} = 1$ for all i . Thus deltas across assets only differ by the scaling variable n_i . The delta for each single asset is not made dependent on the stochastic properties of that asset.

This approach approximates the stochastic properties of the portfolio. An alternative approach, where the value of the portfolio is approximated with a lognormally distributed variable, will be presented in a later section.

3.3.2 The Suggested Model

Olsen and Stensland (1992) studied a similar stopping problem to the *original problem* stated above. They studied optimal timing of real investments when the cost components of a project are additive. In the following we apply the same idea in our setting of downside protection of portfolios. The following decomposition of the *original problem* is central to their idea and constitutes the core of our approximation.

$$\begin{aligned}
 p^*(W(t)) &\equiv \sup_{\tau \in \mathcal{T}(t)} E_t^Q \left[e^{-r\tau} \left(K - \sum_{i=0}^d n_i S_i(\tau) \right)^+ \right] \\
 &= \sup_{\tau \in \mathcal{T}(t)} E_t^Q \left[e^{-r\tau} \left(\sum_{i=0}^d [\lambda_i K - n_i S_i(\tau)] \right)^+ \right] \\
 &\leq \sum_{i=0}^d \sup_{\tau_i \in \mathcal{T}(t)} E_t^Q [e^{-r\tau_i} (\lambda_i K - n_i S_i(\tau_i))^+] \equiv \bar{p}^*(W(t)) \quad (3.3.4)
 \end{aligned}$$

where λ_i is an arbitrary number such that $\sum_{i=0}^d \lambda_i = 1$. The approximation is formulated by allocating a fraction of the strike, λ_i , to each asset in the portfolio. Rather than evaluating one option on the entire portfolio, we evaluate the sum of options on each single asset. The strike for each partial option is an allocated portion of the aggregate strike. The rule for how this allocation should be made, that is, how to choose the λ_i 's, is seen from the inequality in the last line of (3.3.4). The inequality follows from the additional flexibility of being able to select an individual stopping time for each process i , rather than finding a stopping time at which to stop all processes simultaneously. Hence, the approximation represents an upper bound for the value $p^*(W(t))$. Therefore the natural choice of the λ_i 's are those that minimize $\bar{p}^*(W(t))$.

The approximation described above is used to provide a formula for the holdings in a replicating strategy. Because we reconsider the approximating strategy at each re-balancing point in time, the λ_i 's can be re-optimized for each realization of $\mathbf{S}(t)$. Because the optimal λ_i 's varies through time, the suggested procedure is not straightforward to apply to obtain a

closed form pricing formula for the option on the portfolio.

For each i in the last line of (3.3.4) the solution is readily available. By substituting variables in (3.3.3) we see that the value of a live perpetual put on a single asset, $p^*(S_i)$, is

$$p^*(S_i) = \left(\frac{\lambda_i K}{\gamma_i + 1} \right) \left(\frac{(\gamma_i + 1)n_i S_i(t)}{\gamma_i \lambda_i K} \right)^{-\gamma_i},$$

where $\gamma_i = \frac{2r}{\sigma_i^2}$. The optimal exercise boundary is a constant, $C_i = \frac{2r}{2r + \sigma_i^2}$, i.e., it is optimal to exercise the option the first time $\frac{n_i S_i(t)}{\lambda_i K} \leq C_i$.

After inserting this expression into the last line of (3.3.4), we formulate the program from which to find the optimal λ_i 's for each t .

$$\begin{aligned} \min_{\lambda_i} \bar{p}^*(W(t)) &= \min_{\lambda_i} \sum_{i=0}^d \sup_{\tau_i \in \mathcal{T}(t)} \mathbb{E}_t^Q [e^{-r\tau_i} (\lambda_i K - n_i S_i(\tau_i))^+] \\ &= \min_{\lambda_i} \sum_{i=0}^d \left(\frac{\lambda_i K}{\gamma_i + 1} \right) \left(\frac{(\gamma_i + 1)n_i S(t)}{\gamma_i \lambda_i K} \right)^{-\gamma_i} \end{aligned} \quad (3.3.5)$$

$$\text{s.t.} \quad \sum_{i=0}^d \lambda_i = 1$$

The solution is a straightforward application of the Lagrangian method, but parts of the solution have no closed form characterization. However, a solution is available by numerical methods.

Proposition 2. *The solution to the program (3.3.5) is given by*

$$\bar{p}_{min}^*(W(t)) \equiv \sum_{i=0}^d \left(\frac{\lambda_i^*(t) K}{\gamma_i + 1} \right) \left(\frac{(\gamma_i + 1)n_i S(t)}{\gamma_i \lambda_i^*(t) K} \right)^{-\gamma_i} = \sum_{i=1}^d \frac{n_i S_i}{\gamma_i} \left(\frac{\varphi^*}{K} \right)^{\left(1 + \frac{1}{\gamma_i}\right)},$$

and the optimal λ_i is

$$\lambda_i^*(t) = \frac{\left(\frac{\varphi^*}{K} \right)^{\frac{1}{\gamma_i}} (1 + \gamma_i) n_i S_i(t)}{\gamma_i K},$$

where φ^* solves the equation $h(\varphi^*) = 0$, for $h(\varphi) \equiv \sum_{i=1}^d \frac{\left(\frac{\varphi}{K} \right)^{\frac{1}{\gamma_i}} (1 + \gamma_i) n_i S_i(t)}{\gamma_i K} - 1$.

Proof. The Lagrangian for (3.3.5) is

$$\mathcal{L} = \sum_{i=1}^d \left[\left(\frac{\lambda_i K}{\gamma_i + 1} \right) \left(\frac{(\gamma_i + 1)n_i S_i(t)}{\gamma_i \lambda_i K} \right)^{-\gamma_i} \right] - \varphi \left(\sum_{i=1}^d \lambda_i - 1 \right),$$

where φ is the Lagrange multiplier. Let $g(\lambda_i) = \left(\frac{\lambda_i K}{\gamma_i + 1}\right) \left(\frac{(\gamma_i + 1)n_i S_i(t)}{\gamma_i \lambda_i K}\right)^{-\gamma_i}$. Then $g'(\lambda_i) = K \left(\frac{\gamma_i \lambda_i K}{(1 + \gamma_i)n_i S_i(t)}\right)^{\gamma_i}$. $g(\lambda_i)$ is increasing and convex. Hence the $(d + 1)$ first order necessary and sufficient conditions are

$$K \left(\frac{\gamma_i \lambda_i K}{(1 + \gamma_i)n_i S_i(t)}\right)^{\gamma_i} - \varphi = 0 \quad (3.3.6)$$

and

$$\sum_{i=0}^d \lambda_i = 1. \quad (3.3.7)$$

Solving (3.3.6) for λ_i and substituting in (3.3.7) yields the condition $h(\varphi) = 0$. The root of this equation, φ^* , is substituted for φ in (3.3.6), to yield the optimal λ_i . \square

Next, the exercise condition needs to be discussed. Note that an exercise condition for each i is available from the Samuelson (1965) formula. Due to the fact that λ_i is optimized for every t , and because $\bar{p}^*(W(t))$ is a convex function in $S_i(t)$, it will never be optimal to exercise only one option on the basis of this condition. Suppose the opposite, that the exercise condition is satisfied for the option on asset i only. Then λ_i can be reduced so that the exercise condition is not satisfied, which will result in a decrease in $p(S_i(t))$. The corresponding increase in other λ_i 's will increase $\bar{p}^*(W(t))$ by less than this because of the convexity and because these options are less in the money. This implies that such a collection of λ_i 's can not be optimal. Hence, either all or none of the d individual exercise conditions are satisfied.

However, a weaker but still sufficient exercise condition is available. The following proposition is a modified version of proposition 2 in Olsen and Stensland (1992).

Proposition 3 (Olsen and Stensland (1992)). *A sufficient, but not necessary, condition for exercising a perpetual put option on the portfolio $\sum_{i=1}^d n_i S_i$, to yield the payoff $K - \sum_{i=1}^d n_i S_i$, is*

$$K \geq \sum_{i=1}^d \frac{1}{C_i} n_i S_i(t)$$

Proof. Suppose the above condition is satisfied. Define $\lambda_i = \frac{n_i S_i}{C_i K}$ for $i = 1, \dots, d - 1$, and $\lambda_d = 1 - \sum_{i=1}^{d-1} \lambda_i$. From the claimed condition it follows that $\lambda_d \geq \frac{n_d S_d}{C_d K}$. The i 'th term in the sum

$$\bar{p}^*(W(t)) := \sum_{i=0}^d \sup_{\tau_i \in \mathcal{T}(t)} E_t^Q [e^{-r\tau_i} (\lambda_i K - n_i S_i(\tau_i))^+]$$

is now equal to $\lambda_i K - n_i S_i$, i.e., it is optimal to exercise, since $\frac{n_i S_i}{\lambda_i K} \leq C_i$ by definition of λ_i . Accordingly, the entire sum is equal to $K - \sum_{i=0}^d n_i S_i = p(W(t))$. We now see that $p(W(t)) = \bar{p}^*(W(t)) \geq p^*(W(t))$. The last inequality follows from equation (3.3.4). Hence \mathbf{S} is contained in the optimal stopping region for the original problem when $K \geq \sum_{i=1}^d \frac{1}{C_i} n_i S_i(t)$. \square

The perpetual put option can be replicated in a standard way, by managing a portfolio to provide the same sensitivity to the underlying asset as the option. This is achieved by varying the amounts held of the underlying asset and a risk free asset.

Let $Q(t)$ denote the value of the replicating portfolio at time t . If α is the value of the risk free position and θ is the value of the risky position, then

$$Q(t) = \alpha(t) + \theta(t).$$

Applying the standard replication argument, we want $\theta(t)$ to be equal to $\sum_{i=1}^d \delta_i(t) n_i S_i(t)$, where δ_i is the derivative of the value of the option with respect to S_i . In our case we take the derivative of the approximated value $\bar{p}_{min}^*(W(t))$ with respect to S_i and use the proposed exercise condition to indicate when to fully immunize $W(t)$ by letting the short holdings in the replicating portfolio be equal to the long holdings in $W(T)$. Hence,

$$\delta_i(t) = \begin{cases} \frac{\partial \bar{p}_{min}^*(W(t))}{\partial S_i} & \text{if } K < \sum_{i=1}^d \frac{1}{C_i} n_i S_i(t), \\ -1 & \text{if } K \geq \sum_{i=1}^d \frac{1}{C_i} n_i S_i(t), \end{cases}$$

where

$$\frac{\partial \bar{p}_{min}^*(W(t))}{\partial S_i} = -\frac{\gamma_i}{S_i} \left(\frac{\lambda_i^*(t) K}{\gamma_i + 1} \right) \left(\frac{(\gamma_i + 1) n_i S_i(t)}{\gamma_i \lambda_i^*(t) K} \right)^{-\gamma_i}.$$

Note that $\delta_i(t) < 0$. We let $Q(t)$ be zero for $t = 0$ by letting $\alpha(0) = -\sum_{i=1}^d \delta_i(0) n_i S_i(0)$. Hence, we initially have a portfolio of zero value, consisting of short positions in the risky assets and a long position in the risk-free asset. The portfolio is discretely re-balanced to hold $\delta_i(t)$ short of asset i at every re-balancing time. The proceeds are reallocated to or from the risk-free asset. The strategy implies selling the risky assets after a fall in market value and buying back the assets after a market increase, thus accumulating a loss. This loss is the cost of the option or strategy. Termination of the dynamic strategy occurs when the exercise condition is satisfied.

To better understand how the δ_i 's behave as prices change, we study how market conditions influence the determination of λ_i . Remember that lambda serves to distribute the strike of the aggregate option to a strike for each partial option. In the optimization of lambda, lambda is set so that a marginal change in any lambda produces the same change in the value of the

aggregate option, which is the objective function. This is equivalent to letting each option have the same sensitivity with respect to changes in its allocated portion of the strike. In a setting where all assets have equal volatility, this means that the strike for each asset is set so that the intrinsic value of each option constitutes the same fraction of the price of each asset. In other words, each option is made relatively equally into or out of the money by the optimal λ_i . Hence, when all volatilities are equal, all δ_i 's are equal irrespective of the value of each individual asset. Moreover, all deltas change equally with changes in the value of the portfolio.

When volatilities differ, this is not the case. A different volatility will change the sensitivity of an option with respect to changes in the strike. Hence, the optimal lambda for this asset will bring the option sufficiently into or out of the money by adjusting its strike to make that sensitivity equal across all assets. Volatility enters the expression for the δ_i both through λ_i^* and through γ_i . The effect through γ_i is positive, i.e., an increased volatility decreases γ_i , which in turn increases the δ_i for this asset. An increased lambda decreases δ_i , but the effect of increased volatility for one asset on that asset's lambda is not clear. If the aggregate option is sufficiently into the money, i.e., the value of the portfolio is less than the strike, a ceteri paribus change in volatility for one asset may increase the optimal lambda for that asset. Otherwise, when the option is out of the money, less of the strike is allocated to high volatility assets.

However, in total an increased volatility will always increase the delta for that asset. Hence, in the replicating strategy one will always hold a smaller short position in high volatility assets than in low volatility assets. This effect is illustrated in Figure 3.1.

In this figure we see the deltas in a four-asset example of the proposed strategy, denoted *Strategy A*. All assets have different volatilities ranging from 8% to 35%. Each graph shows the delta for one asset when the values of all but the high volatility asset are constant. Lambda is optimized for each value on the x-axis. The values of the assets are 90, 95 and 105, while the fourth asset's value range from 1 to 250. The fourth asset is the most volatile. The strike of the option is 100, hence it is at the money when the fourth asset's value is 110. We see the smooth and equalized nature of how the deltas change in value. The value of 1 for the fourth asset is not sufficiently low to meet the exercise criterion, hence the deltas are not yet -1. Note how the values of the deltas converge.

In Figure 3.2, the same picture is drawn when the fourth asset has a constant value of 110 while the *first* asset range from 1 to 250. The first asset is the asset with the lowest volatility. There is a striking difference in the figures even though the range in the value of the portfolio is the same in both cases. We see that the delta changes a lot less when a low volatility asset




Figure 3.1: A plot of deltas in *Strategy A*. Varying the value of asset S_4 .

changes value than if a high volatility asset moves. In both figures we see that all deltas change in similar proportions, the effect of a change in the value of one asset is absorbed in all the deltas. Furthermore we see that the relative level of the deltas is ranged according to the asset's volatility. The high volatility asset has the highest delta and the low volatility asset has the lowest delta. The deltas are equal in the two figures at 110 on the x-axis in Figure 3.1 and 90 on the x-axis in Figure 3.2.

At the value 1 on the x-axis in the two figures, the average delta is lower in Figure 3.1 than in Figure 3.2. When the high volatility asset has a low price as in Figure 3.1, the volatility of the portfolio is lower than if the low volatility asset has a low price because of the way in which the relative weights change. Hence, we see that the average deltas increase, i.e., becomes closer to zero, when the volatility of the portfolio is high.

In Figure 3.3 the strike in the strategy has been changed from 100 to 600 to produce a graph of the deltas in a range in which the exercise condition is satisfied. We note the convergence of the deltas to -1.




Figure 3.2: A plot of deltas in *Strategy A*. Varying the value of asset S_1 .




Figure 3.3: A plot of deltas in *Strategy A*. Varying the value of asset S_1 .

3.3.3 The Geometrically Weighted Portfolio

In this section we present an alternative way of obtaining a floor for the portfolio. We use a geometrically weighted portfolio, which we know actually follows GBM, to approximate our true portfolio W .

The geometrically weighted portfolio, W_g , we define as

$$W_g(t) = \left(\prod_{i=1}^d S_i(t)^{w_i} \right) \left(\sum_{i=1}^d n_i \right), \quad (3.3.8)$$

The geometric weights, w_i , we define as $w_i \equiv \frac{n_i}{\sum_{i=1}^d n_i}$, where as before n_i is the number of asset i in possession. The assumption of constant n_i is maintained.

As a comparison, the portfolio studied in the preceding section,

$$W(t) = \sum_{i=1}^d n_i S_i(t) = \left(\sum_{i=1}^d \frac{n_i}{\sum_{i=1}^d n_i} S_i(t) \right) \left(\sum_{i=1}^d n_i \right),$$

can be thought of as the arithmetic average asset price in the portfolio times the total number of assets. In (3.3.8) we have defined $W_g(t)$ as the geometric average asset price, times the total number of assets. Because a geometric average is always less than the corresponding arithmetic average, $W_g(t) < W(t)$.

How well W_g approximates W depends among other things on the dispersion in asset prices. If the dispersion is large the approximation is less good than if prices are similar. Another property with this definition is that the same true value of the portfolio, W , may have different approximating values, W_g , depending on the composition of the portfolio. In other applications of the geometric average, e.g., in Asian options, it is natural to use the equally weighted average. The definition above allows different holdings across assets, but simplifies to $\prod_{i=1}^d S_i(t)^{\frac{1}{d}}$ when $n_i = \frac{1}{d}$ for all i . This is the assumption in the numerical simulations that follow later on. A different application of the construction in (3.3.8) can be found in Zhang (1997).

The distribution of W_g is lognormal. Thus, an application of Ito's lemma yields the drift and diffusion terms for dW_g under the risk-adjusted measure Q . Thus

$$dW_g = \mu_g W_g dt + \sigma_g W_g dz^Q,$$

where

$$\begin{aligned} \mu_g &= r - \frac{1}{2} \left(\sum_{i=1}^d (w_i^2 - w_i) \sigma_i^2 - 2 \sum_{i=1}^{d-1} \sum_{j>i}^d w_i w_j \rho_{ij} \sigma_i \sigma_j \right), \\ \sigma_g &= \left(\sum_{i=1}^d w_i^2 \sigma_i^2 + 2 \sum_{i=1}^{d-1} \sum_{j>i}^d w_i w_j \rho_{ij} \sigma_i \sigma_j \right)^{\frac{1}{2}}, \\ z^Q(t) &= \sigma_g^{-1} \sum_{i=1}^d w_i \sigma_i \omega_i^Q(t). \end{aligned}$$

It is easy to verify that $z^Q(t)$ is a Wiener process under Q , noting that $\sigma_i \omega_i^Q(t)$ is the i 'th element of the $d \times 1$ vector of *dependent* Wiener-processes $\mathbf{b}^Q(t) = \mathbf{V} \cdot \mathbf{w}^Q(t)$. Hence, a single asset evolves according to $dS_i = r S_i dt + \sigma_i S_i d\omega_i^Q$, and the correlation coefficient between S_i and S_j is denoted by ρ_{ij} .

With these parameters at hand, an exact value of a perpetual American put option on the geometric portfolio W_g is available simply by replacing r and σ_i with μ_g and σ_g in (3.3.3). Thus

$$p^*(W_g) = \begin{cases} \left(\frac{K_g}{\nu+1} \right) \left(\frac{(\nu+1)W_g(t)}{\nu K_g} \right)^{-\nu} & \text{if } W_g(t) > \frac{2\mu_g}{2\mu_g + \sigma_g^2} K_g, \\ K_g - W_g(t) & \text{if } W_g(t) \leq \frac{2\mu_g}{2\mu_g + \sigma_g^2} K_g, \end{cases}$$

where $\nu = \frac{2\mu_g}{\sigma_g^2}$.

Because $W_g(t) < W(t)$, it follows that $p^*(W_g(t)) > p^*(W(t))$ if $K_g = K$, as the value of a put option increases when the value of the underlying assets decrease. Thus, both $p^*(W_g(t))$

and $\bar{p}_{min}^*(W(t))$ approximate $p^*(W(t))$ from above. However, while an exact replicating strategy can be found in the former case for an option on the approximated value of $W(t)$, in the latter case an approximated strategy has been found for an option on the exact value of $W(t)$.

A replicating strategy can be formulated in a similar fashion as in the former section. Let $Q_g(t)$ denote the value of the replicating portfolio at time t . As before we have that

$$Q_g(t) = \alpha_g(t) + \theta_g(t)$$

where subscript g indicates that the holding is for the replicating portfolio of the geometrically weighted portfolio. Again, let $\theta_g(t) = \sum_{i=1}^d \delta_{g,i}(t)n_i S_i(t)$. $\delta_{g,i}$ is the delta of the option, $p^*(W_g)$, with respect to asset i , but with some adjustments. Note that

$$\frac{\partial p^*(W_g)}{\partial S_i} = \frac{dp^*(W_g)}{dW_g} \frac{\partial W_g}{\partial S_i} = \begin{cases} -\frac{\nu w_i}{S_i} p^*(W_g) & \text{if } W_g(t) > \frac{2\mu_g}{2\mu_g + \sigma_g^2} K_g. \\ -\frac{w_i}{S_i} W_g & \text{if } W_g(t) \leq \frac{2\mu_g}{2\mu_g + \sigma_g^2} K_g. \end{cases}$$

There is nothing here to prevent this derivative from becoming less than -1 . In the context of a replicating strategy this implies that the desired short position may become larger than the long position that is being protected. Outright short positions are at odds with many portfolio guidelines. Moreover, outright short positions are implicitly ruled out in the strategy of the former section with which we want to compare this strategy. Hence, to make a fair comparison we put a lower limit of -1 on the deltas.

Even though $\frac{dp^*(W_g)}{dW_g(t)}$ converges to -1 when $W_g(t)$ approaches the exercise boundary, $\frac{\partial p^*(W_g)}{\partial S_i}$ does not. This is because of the multiplicative rather than additive nature of $W_g(t)$. Recall that we are trying to replicate a floor for the value of $W(t)$ by capturing the exposure of an option on $W_g(t)$. When the exercise criterion is met, following this strategy will produce a portfolio that immunizes $W_g(t)$. What we want is rather a portfolio that immunizes $W(t)$, i.e., a hedge ratio of -1 for each asset when the exercise criterion is satisfied.

Hence, we let the replicating positions in the risky assets be

$$\delta_{g,i} = \begin{cases} \max(-\frac{\nu w_i}{S_i} p^*(W_g), -1) & \text{if } W_g(t) > \frac{2\mu_g}{2\mu_g + \sigma_g^2} K_g. \\ -1 & \text{if } W_g(t) \leq \frac{2\mu_g}{2\mu_g + \sigma_g^2} K_g. \end{cases}$$

In other respect the replicating strategy is conducted in the same fashion as in the former case. We let $Q_g(0)$ equal zero by letting $\alpha_g(0) = -\sum_i^d \delta_{g,i}(0)n_i S_i(0)$, and re-balance the holdings to keep the risky position in each asset equal to $\delta_{g,i}$ at each re-balancing time.

Note that the only element in $\theta_{g,i}$ that varies across assets is S_i . The strategy does not discriminate on volatility as *Strategy A* did. This means that two assets with different volatilities but similar values, will enter the replicating strategy with the same short holdings. The




Figure 3.4: A plot of deltas in *Strategy B*. Varying the value of asset S_4 .

volatility of each asset is averaged to a single measure for $W_g(t)$ through the computation of σ_g .

In Figure 3.4 we show the same plot as in Figure 3.1 but with deltas computed for the strategy of this section, *Strategy B*. First, note the constant differences between the deltas for assets 1 through 3. If the values of these assets had been equal there would have been no difference between the deltas. The levels of the deltas are ranged according to the values of the assets. The changes in these three deltas are only a function of the changing value of $W_g(t)$ as the value of the fourth asset changes.

Note also the binding constraint of -1 for the delta of the fourth asset and the non-convergence of the other deltas. We see that the exercise constraint in this strategy is binding whereas it was not in Figure 3.1. Note that because some $\theta_{g,i}$ may be larger than -1 at exercise, a large change may take place in the position of such an asset. This makes *Strategy B* more vulnerable to large price movements and adds more operational risk to the strategy.

In Figure 3.5 the setting is as in Figure 3.2. The striking difference between *Strategy A* and *Strategy B* is well illustrated by comparing Figure 3.5 and Figure 3.2. Whereas in *Strategy A* it made a big difference whether a low or a high volatility asset changed in value, it has no impact in *Strategy B*. The only difference between Figure 3.4 and Figure 3.5 is due to the




Figure 3.5: A plot of deltas in *Strategy B*. Varying the value of asset S_1 .

difference in the constant values of the first and the fourth asset, i.e., in Figure 3.4 S_1 is held constant at 90 while in Figure 3.5 S_4 is held constant at 110.

3.4 Numerical Testing

In this section we compare the performance of the two strategies by means of a Monte Carlo simulation of a market containing four risky assets and one risk-free asset. That number of assets in the simulation is chosen arbitrarily.

Monte Carlo simulations are often used to numerically compute the value of a derivative. The price process for the underlying asset must be generated under the risk-adjusted measure Q , and a sample discounted expected payoff from the derivative is computed. In our case simulations are used to test replication strategies that are supposed to be conducted with observed prices described by the probability measure P .

The costs related to the proposed replication strategies are not dependent on the expected returns of the portfolio. However, if the individual assets in the portfolio are simulated with different trend parameters, the comparison of the strategies will be affected. This is easy to see noting that the strategies differ in the amounts held short in the different assets, i.e., the

deltas.

To avoid letting the analysis be contaminated by some arbitrarily chosen set of expected returns, we let the expected return, i.e., the trend parameter in the simulation, be equal across all assets.

In this chapter we are studying time-invariant strategies; time does not enter into the decision of how to compose the replicating portfolios. However, in order to make our test end at some point, we need to choose a time horizon on which the payoffs are to be compared. The two protective strategies might be exercised at different times since they have different exercise criteria. By “exercise” is meant a situation where the strategy deltas are all -1, after which the portfolio is held in the risk-free asset until the time horizon. The values at the horizon are discounted back to $t = 0$ with the risk-free interest rate to facilitate comparison. The time horizon is set arbitrarily to a long period such as 10 years.

The payoff from the buy-and-hold strategy is simply the value of a portfolio containing one of each risky assets, $W(T)$, at the horizon T . There is no trade in the period studied.

The end value of the protected strategies, A and B , is the sum of the end value of the buy and hold strategy and the end value of the replicating portfolio, i.e., $A = W(T) + Q(T)$ and $B = W(T) + Q_g(T)$. As noted earlier $Q(0) = Q_g(0) = 0$. The risky component of $Q(t)$ and $Q_g(t)$ is always equal to the current set of deltas multiplied by the current set of prices. The risk-free component is the former period's risk-free component with one period of interest minus the necessary trade to move from last period's deltas to the current deltas. Hence, the value of the risk-free position, $\alpha(t)$, evolves according to

$$\alpha(t) = \alpha(t_{-1})e^{r(t-t_{-1})} - \sum_{i=1}^d (\delta_i(t) - \delta_i(t_{-1})) S_i(t),$$

and the value of the risky position is $\sum_{i=1}^d \delta_i(t) S_i(t)$. Here, t_{-1} denotes one period earlier. The sum of the risky and the risk-free position is the accumulated profit or loss of the replicating portfolio since the value at $t = 0$ is zero. The only difference in $Q(t)$ and $Q_g(t)$ is in the way the deltas are calculated, as described earlier.

In the following simulation we study a market of one risk-free and four risky assets. The risky assets have the correlation matrix given in Table 3.1, where the initial values of each asset are given in the last line of the table. The volatilities are given along the diagonal.

Price paths for each asset are simulated by the following expression:

$$\mathbf{S}(t) = \mathbf{S}(t_{-1}) \cdot \exp \left[\left(r\mathbf{1} - \frac{1}{2} \mathbf{v}^2 \right) dt + \sqrt{dt} \cdot \mathbf{V} \cdot \mathbf{Z} \right].$$

Table 3.1: Correlation structure.

Asset	S_1	S_2	S_3	S_4
S_1	0.080			
S_2	0.375	0.150		
S_3	0.010	0.100	0.250	
S_4	0.020	0.100	0.000	0.350
$S_i(0)$	90	95	105	110

Here \mathbf{S} is a 4×1 vector of asset prices, \mathbf{v} a 4×1 vector of volatility, \mathbf{V} a 4×4 matrix such that $\mathbf{V} \cdot \mathbf{V}^T$ is the covariance matrix developed from 3.1, and \mathbf{Z} is a 4×1 vector of standard normal variates. The notation ".*" and ".2" means that the operation should be performed element by element.

The buy and hold portfolio consists of holdings of $\frac{1}{4}$ of each asset, so that $W(0) = 100$. The strike, K , of the two strategies is set at 100 as well, hence we want to protect the nominal value of the portfolio. We chose the trend parameter to be equal to the risk-free interest rate, which is 4%. The replicating portfolio is re-balanced daily and we compare the payoffs after ten years and at the time one of the strategies is exercised.

Table 3.2 shows some diagnostics of the simulated prices. In the table, sample statistics from the simulation are compared with their true values as given by the parameters of the simulation. The notation $\hat{E}[\cdot]$ means empirical expectation. From these numbers it looks as

Table 3.2: Descriptive statistics.

Asset	$r - \frac{1}{2}\sigma_i^2$	$\hat{E}[\ln \frac{S_i(T)}{S_i(0)}]$	$S_i(0)e^{rT}$	$\hat{E}[S_i(T)]$
S_1	3.78%	3.72%	134.26	134.94
S_2	2.88%	2.83%	141.72	141.16
S_3	0.88%	0.94%	156.64	156.98
S_4	-2.13%	-2.14%	164.10	163.52

though the simulation is reasonably unbiased. Here we have simulated 10000 price paths.

Table 3.3 shows some summary statistics for the two strategies. To compare the results of the simulation with strategies with known answers, the cost-of-hedging statistics for replicating portfolios of a 10-year European put on one of the assets, S_3 , and of a perpetual put on the same single asset are included. The replicating portfolios are constructed in the same fashion in all cases. Only the computation of the deltas differs. Again, the same 10000 simulations of

10-year long, daily, price paths are used.

The cost of hedging comes from the average difference between the price paid for the asset and the price realized for it as the replicating portfolio is dynamically re-balanced. In the case of a European put, the cost is the discounted difference between the risk-free holdings and the strike when the option ends in the money, and the discounted sum of the risky and the risk-free holdings when the option ends out of the money. However, since the delta of a European option converges to zero when the option ends out of the money, there will always be zero risky holdings in the replicating portfolio at an out of the money expiration. Letting the risk-free position be denoted by α and the risky by θ , we let $cost = (\alpha(T) - K)e^{-rT}$ when $S_3(T) < K$ and $cost = \alpha(T)e^{-rT}$ when $S_3(T) \geq K$. T denotes the horizon, i.e., 10 years. The average cost-of-hedging measured in this way will approximate the Black and Scholes (1973) value of the option.

Because a perpetual option has no horizon, there may be risky holdings in the replicating portfolio even if the strategy is terminated when the option is out of the money at T . The measure of the cost of hedging used for the perpetual strategies is analogous to the measure defined above. In cases where the exercise condition has not been satisfied prior to the (premature) ten-year horizon, the cost is calculated as the present value of the sum of the risky and the risk free positions, i.e., $cost = (\alpha(T) + \theta(T))e^{-rT}$. When the exercise condition has been satisfied prior to the horizon, T , we have held a risk-free position ever since that time. The cost measure is then calculated as the discounted value of the difference between the risk-free position at exercise and strike. That is, $cost = (\alpha(t^*) - K)e^{-r(t^*)}$, where t^* is the time the exercise condition has been satisfied.

In Table 3.3 average cost measures together with their standard errors are given. Theoretical values of the ten-year European and the perpetual put on the single asset are given in the *Cost* column. *StrategyA* is the strategy proposed in this chapter while *StrategyB* is the one based on the geometrically weighted portfolio. Yearly turnover in the replicating strategies is denoted by Φ . Even though all the expected cost measures are below their true values, we see

Table 3.3: Simulated costs.

	$Cost$	$E[cost]$	$\hat{\sigma}(cost)$	$\hat{E}[\Phi]$	$\hat{\sigma}(\Phi)$
<i>European Put</i>	-11.3114	-11.3124	0.3091	151.05	6.44
<i>Perpetual Put</i>	-19.6795	-14.5076	5.9913	202.46	72.70
<i>StrategyA</i>	-	-2.4500	5.8588	124.72	36.55
<i>StrategyB</i>	-	-4.8910	4.3320	182.69	66.14

that *StrategyA* has a lower cost measure than *StrategyB*. These two measures are not comparable with the one for the *Perpetual Put* strategy because this applied to asset S_3 which has a higher initial value and higher volatility than the portfolio underlying the two other strategies. By way of comparison, the theoretical value of a perpetual put with a volatility equal to the initial volatility of the buy-and-hold portfolio (which is 0.1299), a strike of 100 and an initial asset value of 100, is 7.0299.

The Black and Scholes (1973) value of a 10-year European put on an asset with an initial value of 100, a strike of 100 and a volatility of 0.1299, is 2.9061. The numerically computed value of a 10-year European put on $W(t)$ with strike 100, that is, $e^{-rT} \hat{E}^Q[(K - W(T))^+]$, is 2.5344. This estimate is based on the same simulation as above. Because of the unknown distribution of $W(T)$, the value of 2.5344 is the best estimate of a low limit for the cost of the time-invariant strategies, which has greater promise than that of a European put. Hence, it seems that the measure of the expected cost in Table 3.3 has a negative bias. However, $\hat{E}[\text{cost}]$ for the time-invariant strategies are *ad hoc* measures and are not expected to approximate the true cost of the perpetual options. This is because the measure is calculated at a time T which is randomly chosen in the setting of an infinite horizon for the strategies.

In the simulation, *StrategyA* satisfies the exercise criterion 991 times while *StrategyB* is exercised 3160 times out of the 10000 simulations. *StrategyB* is always exercised earlier than *StrategyA*.

The most striking observation in Table 3.3 is the large difference between the expected turnovers of the two strategies. *StrategyB* has an almost 50% higher turnover than *StrategyA*. "Turnover" or Φ , is computed in the following way: The absolute value of each re-balancing trade is summed across all assets and accumulated (without interest) over time. This accumulated sum is then divided by the number of years the strategy is active, i.e., until an exercise criterion is satisfied. There is no turnover or trade in any strategy once it has been exercised. The average over all simulations is then computed. The numbers should be compared with the value of the initial portfolio, 100. Hence *StrategyA* is turned over 1.2 times a year. The relatively low turnover in *StrategyA* is important if there are transaction costs, and a strong argument in favor of this strategy if the performance is not too unfavorable otherwise.

To further assess the performance of the two strategies, we compare the distributions of the discounted end values at T . As before we define the end value of *StrategyA* as $A(T)$ and of *StrategyB* as $B(T)$. The end value of the buy-and-hold strategy is of course $W(T)$. Because all assets have the same expected end-value or growth rate, all strategies should have the same

expected end value. Remember that the initial value of the replicating portfolio is zero and that the portfolio is self-financing in the sense that no funds are withdrawn or invested in this portfolio. Hence the expected end value of the replicating portfolio is zero. The standard deviation of the end-value is expected to decrease with the proportion of the strategy allocated to the risk-free asset. As previously we compute the same measures for the replicating portfolio of a 10-year European put by way of comparison.

In Table 3.4, \mathcal{V} denotes end values discounted to $t = 0$. \mathcal{N} is the number of times the end value of “the underlying” is below the floor and η is the proportion of these cases where the end values of the strategies are below this floor. So η is a measure of how well the floor is holding. ζ denotes the average underscore when the end value of the strategy is below the floor. The floor is defined as the discounted value of the strike less the cost of the option, so for the European put the floor is $100e^{-rT} - 11.3114$. For the two perpetual strategies we calculate the floor with an arbitrarily chosen cost of 3.00 for both strategies, hence the floor is $64.81 = 100e^{-rT} - 3.0000$. In Table 3.4 we recognize approximations for initial value of the

Table 3.4: Performance.

	$E[\mathcal{V}]$	$\hat{\sigma}(\mathcal{V})$	\mathcal{N}	η	ζ
$S_3(T)$	105.2283	97.19	—	—	—
<i>European Put</i>	105.1085	89.84	4244	50.78%	0.43%
$W(T)$	99.9790	52.10	—	—	—
$A(T)$	100.0550	47.19	2228	10.64%	1.94%
$B(T)$	99.9248	45.96	2228	9.74%	2.96%

assets, 105 for $S_3(0)$ and 100 for the other strategies. We also see the higher volatility of the single asset than that of the portfolio W . This explain why the number of in-the-money end values are much higher for the single asset. Both *StrategyA* and *StrategyB* have $W(T)$ as their underlying value, hence the same number of in the money paths. Both strategies hold their floor very well. The under shooting is less frequent but of a greater magnitude than what might be expected when a standard European put option is replicated. Again, *StrategyA* compares favorably relative to *StrategyB*. The *ad hoc* assumption of 3.00 as the cost is a low estimate. Increasing this estimate lowers the floor and the strategies perform even better with respect to not breaking the floor.

In Figure 3.6 we show histograms of $e^{-rT}W(T)$, $e^{-rT}A(T)$ and $e^{-rT}B(T)$.

The end values are discounted to the present and counted in bins of 2, ranging from 0 to 300. The peak at the right in the figure contains all realizations above 300. We see the slightly

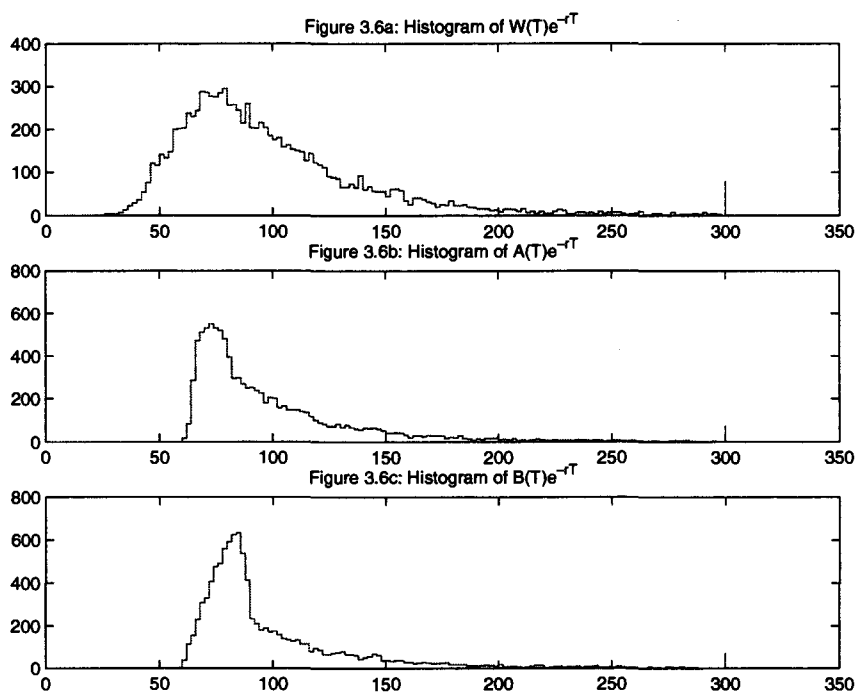


Figure 3.6: Histograms of $e^{-rT}W(T)$, $e^{-rT}A(T)$ and $e^{-rT}B(T)$.




Figure 3.7: Accumulated histograms for W , *Strategy A* and *Strategy B*.

better floor-holding property of *Strategy A* and a favorable performance of *Strategy B* when the end value is in the region 70 to 100. For end values above this level, *Strategy A* has a slight edge. Remember that the expectations are equal in all these three distributions.

In Figure 3.7 the same comparison is presented with accumulated histograms. Hence, from the curve one can read off the proportion of outcomes that is less than the corresponding value on the x-axis. The vertical line indicates the floor of the two competing strategies.

3.5 Concluding Remarks

In this chapter we have proposed a new method for designing a replicating strategy for an option on a buy-and-hold portfolio. The method incorporates the volatility of each individual asset when the delta for that asset is calculated. The method has been outlined here in the context of replicating a Perpetual put option on a portfolio. Furthermore, the strategy has been tested with Monte Carlo simulations and compared with a more standard way of conducting such a replication. The proposed method compares favorably. In particular, the turnover associated with the strategy is reduced by approximately 35% without this making the performance of the strategy poorer.

In view of the promising results, further studies are tempting. A better assessment of the

difference between the proposed method and standard methods might be achieved by studying a more conventional type of option on a portfolio. In our setting it is difficult to provide a good estimate of the true cost of the strategy because of the intangible nature of a Perpetual option.

The correlation between the assets does not explicitly enter into the calculation of *Strategy A*. Implicitly, it effects the probability of whether the exercise criterion is satisfied or not. However, one would think that the method could be further enhanced by finding a way of including the correlation in the calculation of the deltas. This is a topic for further studies.

The method may be further investigated by using it on other types of options. In particular, it would be interesting to see how the strategy performs when the strike is the value of a portfolio as well. That is, it would be interesting to test a strategy of replicating an option to sell the best performing of two portfolios, or in other words, to provide downside protection of the difference between a portfolio and its benchmark portfolio. Such a setting might be particularly relevant for tactical asset allocation strategies which typically involve asset classes with very different volatilities.

Chapter 4

Options on Active Portfolios: Pricing Under Semi-Lognormality

An active portfolio is a difference in holdings between two portfolios. This chapter analyzes approximation formulas for the prices of European options on such long-short portfolios. An exact pricing formula is not available in a Black and Scholes world because the distribution of sums of lognormal variables is unknown. We test how well an assumption of semi-lognormality, that is, a joint lognormal-normal distribution of prices and differences between portfolios, performs in pricing formulas. Numerical tests show discouraging results relative to a standard approximation. Along the way, pricing formulas for some rather exotic derivatives are developed as preliminary results.

4.1 Introduction

A standard setting in the money management industry is the following: A sponsor defines a benchmark portfolio representing his long term risk-reward tradeoff. Often this is chosen to be a market capitalization weighted index representing an average investor's holdings. An active manager is allowed to deviate from the benchmark portfolio within limits, and his fees are often tied to the degree of out-performance relative to the benchmark. An active manager will deviate from the benchmark, i.e., form an active portfolio, if he disagrees with the views or expectations that make the benchmark an efficient portfolio. The manager is said to disagree with consensus, given that the benchmark is the average market portfolio. To deliver returns in excess of the benchmark, he must possess predictive powers.

In this chapter we study pricing formulas for European options on an active portfolio, i.e., on the difference between a portfolio and a benchmark portfolio. Such a difference portfolio is commonly termed a long-short portfolio. We focus on portfolio exchange options, which yield

the maximum of the difference between two portfolios and zero. We also briefly comment on portfolio spread options which yield the maximum of the difference between two portfolios and a constant, i.e. have a strike different from zero.

There are at least two arguments for why such an option on the active portfolio should interest a manager. The first is related to the common practice of using stop-loss policies. A stop-loss policy is a predetermined set of variables, most often prices, which guide the manager as to when to unwind his active portfolio, or in other words when to return to the benchmark portfolio. There are many different kinds of motivations for such a policy. Hypotheses testing could be one, where the variables are set in a way not likely to be observed given the expectations the active position is based on. Hence, if the specified set of variables should indeed be observed, the rationale for the active portfolio is probably wrong, and the position should be canceled. However, the argument could just as well be used to get rid of an option on the active portfolio as to get rid of the active portfolio itself.

Another motivation could simply be a strong aversion against losing more than the amount allowed for in the stop-loss policy. This is usually connected to some accounting period, i.e., a horizon, for which the accumulated result should look good. In this case one could argue that an option would be a more precise instrument. The reason lies in the difference between a stop-loss policy and a (call) option on the active portfolio. A stop-loss policy produces a path-dependent payoff. The position could be stopped out at an early stage so that the investor gives up a favorable price change from that point forward, prior to the fixed horizon. A European option, on the other hand, is not path-dependent. Regardless of how the price moves from now until the horizon, the payoff will be the same: a function of the prices at the horizon. However, this certainty does not come without the cost of the option.

The second argument is more theoretical, and presented in Leland (1980). In that article it is discussed how investors who think they benefit from investing in options or portfolio insurance, i.e., a convex payoff schedule, differs from a representative investor in the market. One result is that investors with average risk tolerance, average opinion of riskiness, but a more optimistic assessment of the expected return than the average, would have greater expected utility from a convex payoff schedule than from a buy and hold strategy. An intuitive explanation of this result could be that investing in a convex payoff makes it possible to better exploit a superior ability to spot good investments, while at the same time keeping the risk within tolerable levels.

By definition, an active manager holds the difference between his portfolio and a benchmark

portfolio because he believes that combination of assets to be undervalued. Thus, following Leland (1980) his expected utility would increase by holding a contingent claim yielding a convex payoff from that difference portfolio rather than buying and holding the difference portfolio outright.

This chapter is only concerned with convex payoffs of a certain type: European options on the active portfolio. The reason why pricing of derivatives involving portfolios is not straightforward, is that the probability distribution of a portfolio, being a sum of assets, is unknown when each asset price is lognormally distributed. Long-short portfolios in particular are certainly not lognormally distributed, since they are able to take on negative values. The standard way of finding approximating values for such options is nevertheless to treat the portfolios *as if* they were lognormally distributed.

The aim of this chapter is to develop and test several pricing approximations. The formulas are tested by comparing the approximated price with the true price as given by Monte Carlo simulation, that is, a true price given an assumption about the probability distribution of the individual assets.

The main result in this chapter stems from a way of rewriting the payoff from the option in terms of *payout assets* and *trigger assets*. The payout assets determine the size of the payoff, while the trigger assets determine whether there is going to be any payoff at all. It is then possible to assume a different probability distribution for each class of assets. We develop and test a pricing formula in which all assets are joint lognormally distributed, differences between two portfolios are normally distributed and the joint distribution of assets and portfolio differences is lognormal-normal, or semi-lognormal. This distribution is described in Crow and Shimizu (1988).

This pricing formula provides an explicit link between individual asset prices, their individual risk parameters and the value of the portfolio exchange option. This yields some insight relative to the standard approximation, which is only a function of the risk characteristics of the aggregated portfolios. However, the Monte Carlo testing shows that pricing accuracy is not improved compared with the standard approach.

The chapter is organized in the following way: Section 2 outlines the basic features of the financial model we use. Section 3 describes the standard method used to approximate the value of portfolio exchange options. In Section 4, the method of externalizing the trigger assets is described, and a few preliminary results are given which are necessary to permit the use of the approach under semi-lognormality. This assumption, i.e., a joint lognormal-normal

or semi-lognormal distribution, is introduced in Section 5, where the formula for a portfolio spread option is also given. In Section 6, the various pricing relations are numerically tested and compared with the "true value" as given by Monte Carlo simulation. Section 7 is the conclusion.

4.2 Financial Model

We adopt a standard financial market model, often termed a generalized Black and Scholes model. The model is generalized in the sense that there are N risky assets rather than one. The price of risky assets, $S_i(t)$ follows an N -dimensional geometric Brownian motion. Drift and volatility are assumed to be constant. Hence, in a complete probability space we let

$$dS_i(t) = S_i(t) \left[\mu_i dt + \sum_{j=1}^N \hat{\sigma}_{ij} dW_j^{\mathbb{P}}(t) \right], \quad (4.2.1)$$

where μ_i is the expected growth rate, $\hat{\sigma} = \{\hat{\sigma}_{ij}\}$ is an $N \times N$ matrix such that $\hat{\sigma} \hat{\sigma}^\top = \sigma$ is the covariance matrix, and $W^{\mathbb{P}}(t)$ is a Wiener process under the subjective probability measure \mathbb{P} . Asset number $N + 1$ is a risk free-asset with constant return of r .

Because absence of arbitrage and dynamical completeness is assumed, there exists an unique equivalent martingale measure, or yield-equating measure, \mathbb{Q}_1 given by

$$\frac{d\mathbb{Q}_1}{d\mathbb{P}} = \exp \left[-\frac{1}{2} \|\theta\|^2 dt - \theta^\top dW^{\mathbb{P}}(t) \right], \quad (4.2.2)$$

where θ is the N -vector of risk premiums, i.e., $\sigma\theta = \mu - r\mathbf{1}$. Under \mathbb{Q}_1 all assets follows a process as in (4.2.1), but with the drift parameter μ_i replaced by r . A clear and thorough exposition of the model is given in Karatzas (1997).

4.3 Standard Approximation of European Portfolio Exchange Options

A *spread option* is an option on the difference, i.e., the spread, between two assets. The name *exchange option* is also used, but mostly in cases where the strike of the spread option is zero. In that case the holder receives one asset in exchange for the other asset, i.e., pays zero for the spread between the assets. We study a case in which the assets involved are portfolios, and we focus mainly on a zero strike.

The standard approach when pricing portfolio exchange options (see Zhang (1997)), is to approximate the two portfolio values using a geometric mean, compute the parameters of the geometric process of this approximation, and pretend that the arithmetic portfolio values follow this process. The result is a pricing formula that treats the portfolio values as if they were single assets.

Under these assumptions, the value of a portfolio exchange option at time t , with payoff at time T equal to

$$P(I_1, I_2, T) = (I_1(T) - I_2(T))^+, \quad (4.3.1)$$

and assuming that the portfolio values I_1 and I_2 follow a geometric Brownian motion, is

$$P(I_1, I_2, t) = I_1(t)N(d_1) - I_2(t)N(d_2)$$

where

$$d_1 = \frac{\ln \frac{I_1}{I_2} + \frac{\tau}{2}\sigma^2}{\sigma\sqrt{\tau}},$$

$$d_2 = d_1 - \sigma\sqrt{\tau},$$

$$\tau = T - t,$$

and

$$\sigma^2 = \sigma_1^2 - 2\rho_{12}\sigma_1\sigma_2 + \sigma_2^2. \quad (4.3.2)$$

Here σ_i is the volatility of the portfolios, and ρ_{ab} is the correlation between the two portfolios. These parameters change over time when the portfolio holdings are fixed and, moreover, the real distribution of the portfolios is unknown.

A common approximation is achieved by computing the volatilities of the portfolios in the following way:

$$\hat{\sigma}_1^2 = \sum_{i=1}^k \sum_{j=1}^k w_{i1}w_{j1}\sigma_{ij}$$

and

$$\hat{\sigma}_2^2 = \sum_{i=k+1}^n \sum_{j=k+1}^n w_{i2}w_{j2}\sigma_{ij},$$

where σ_{ij} is the variance when $i = j$, or otherwise the covariance between assets i and j . w_{ij} is the relative weight of asset i in portfolio j . The covariance between I_1 and I_2 is approximated with

$$\hat{\sigma}_{12} = \sum_{i=1}^k \sum_{j=k+1}^n w_{i1}w_{j2}\sigma_{ij}.$$

Table 4.1: Base Case parameters.

Asset	$S_i(0)$	S_1	S_2	S_3	S_4
S_1	100	0.10			
S_2	100	0.2	0.15		
S_3	100	0.2	0.3	0.20	
S_4	100	0.1	0.4	0.2	0.25

Table 4.2: Standard approximation in Base Case.

	Monte Carlo	Standard app.	bias
Price	12.8515	12.6025	-1.94%

These parameters are used as proxies for the parameters in the valuation formula (4.3.2).

Note that according to Ito's lemma, the parameters of the process of a geometric portfolio with unitary coefficients, i.e., $\hat{I}_1 = k \left(\prod_{i=1}^k S_i \right)^{\frac{1}{k}}$ are given by the above formulas if w_{i1} is replaced by $\frac{1}{k}$.

In Table 4.1, some parameters of a base case example are given. We let $k = 2$ and $n = 4$. The table displays a correlation matrix for the four assets and their values at time zero. Volatilities are given at the diagonal in the matrix.

In Table 4.2 the approximated price obtained using the standard approach of (4.3.2) is compared with the true value of the option, calculated by means of Monte Carlo simulation. The calculations are based on the setting in Table 4.1. We see that the standard approach has a negative bias of almost 2% in this case. Even though the size of the bias varies for different sets of parameters, it is always negative in reasonable cases. More numerical results are presented in a later section.

In the following section we develop alternative approximation formulas and later we test whether they provide a higher accuracy than the standard approach.

4.4 Externalizing the Trigger Assets

In an ordinary exchange option, the two single assets involved have a double purpose. The difference between the two assets determines the size of the payoff but also determines whether there is going to be any payoff at all. However, there are several examples of exotic options where these two features of the option are separated. To illustrate this point, the payoff from

a standard exchange option may be written as

$$(S_i - S_j)^+ = (S_i - S_j) \mathcal{X}_{(S_i > S_j)},$$

where $\mathcal{X}(\cdot)$ denotes an indicator function with value 1 if the condition is true, and zero otherwise. The arguments of the indicator function are the *trigger assets*, while the assets in the main parenthesis are the *payout assets*. These do not necessarily have to be the same assets, nor do they have to have the same distribution.

In the following we provide a decomposition of the payoff in (4.3.1) that enables us to develop pricing formulas in our portfolio context, where the trigger assets and the payout assets are separated. Hence, we let

$$\begin{aligned} P(\mathbf{S}, T) &= (I_a(T) - I_b(T))^+ \\ &= \left(\sum_{i=1}^k S_i - \sum_{j=k+1}^n S_j \right)^+ \\ &= \left(\sum_{i=1}^k S_i - \sum_{j=k+1}^n S_j \right) \mathcal{X}_{(\sum_{i=1}^k S_i > \sum_{j=k+1}^n S_j)} \\ &= \sum_{i=1}^k \sum_{j=k+1}^n \left[\left(\frac{1}{n-k} S_i - \frac{1}{k} S_j \right) \mathcal{X}_{(I_1(T) > I_2(T))} \right] \end{aligned} \quad (4.4.1)$$

This is the sum of several contingent claims, with the trigger assets common to all of them. The payout assets however, differ for each claim. By the law of one price, a sum of values must equal the value of a sum. Hence, we can find the value of the portfolio exchange option, i.e., $P(\mathbf{S}, t)$, if we can find the value of each bracketed expression in (4.4.1).

To proceed, we again need to assume the stochastic properties of the portfolios, which of course are still unknown. However, with this formulation we express the value of the option in terms of the properties of the individual securities as well, not only in terms of the portfolio properties.

In the next section we develop the pricing relation for the bracketed expression in (4.4.1) when the four assets follow a four-dimensional geometric Brownian motion. In a later section we find approximating pricing relations for the portfolio exchange option under two different assumptions regarding the distribution of the trigger assets, i.e., the portfolios. However, the individual assets follow a multi-dimensional geometric Brownian motion in both cases.

4.4.1 Pricing a Four Asset Exchange Option

We give the name *four-asset exchange option* to a derivative with payoff at T equal to

$$P(\mathbf{S}, T) = (S_1 - S_2)\mathcal{X}_{(S_3 > S_4)}.$$

It is termed so because the owner receives the spread between two assets, given that the third asset is larger than the fourth.

In our model we know that

$$P(\mathbf{S}, t) = E_t^{\mathbb{Q}_1} \left[e^{-r(T-t)} (S_1 - S_2)\mathcal{X}_{(S_3 > S_4)} \right], \quad (4.4.2)$$

where \mathbb{Q}_1 denotes the risk-neutral or yield-equating probability measure. Under this measure prices evolves according to

$$dS_i = S_i r dt + S_i \sigma_i dW_i,$$

where r is the discounting rate. Define $S_b = \frac{S_3}{S_4}$. From Ito's lemma we know that S_b follows

$$dS_b = S_b \left(\bar{\mu}_b dt + \bar{\sigma}_b dW^{\mathbb{Q}_1} \right),$$

where

$$\begin{aligned} \bar{\mu}_b &= \sigma_4^2 - \rho_{34}\sigma_3\sigma_4, \\ \bar{\sigma}_b &= (\sigma_3^2 - 2\rho_{34}\sigma_3\sigma_4 + \sigma_4^2)^{\frac{1}{2}}, \\ W^{\mathbb{Q}_1}(t) &= \sigma_b^{-1} \left(\sigma_3 W_3^{\mathbb{Q}_1}(t) - \sigma_4 W_4^{\mathbb{Q}_1}(t) \right), \end{aligned}$$

and $W^{\mathbb{Q}_1}(t)$ is a Brownian motion under \mathbb{Q}_1 . Then we can rewrite (4.4.2) as follows:

$$P(\mathbf{S}, t) = S_2(t) E_t^{\mathbb{Q}_1} \left[\frac{e^{-r(T-t)} S_2(T)}{S_2(t)} \left(\frac{S_1(T)}{S_2(T)} - 1 \right) \mathcal{X}_{(S_b > 1)} \right] \quad (4.4.3)$$

$$= S_2(t) E_t^{\mathbb{Q}_2} \left[(S_a - 1) \mathcal{X}_{(S_b > 1)} \right], \quad (4.4.4)$$

where $S_a = \frac{S_1}{S_2}$. The last equality follows because $\frac{e^{-r(T-t)} S_2(T)}{S_2(t)}$ may be interpreted as a Radon-Nikodym derivative. Thus we have

$$\frac{d\mathbb{Q}_2}{d\mathbb{Q}_1} = \frac{e^{-r(T-t)} S_2(T)}{S_2(t)} = \exp \left\{ - \left(\frac{1}{2} \sum_{k=1}^4 \hat{\sigma}_{2k}^2 (T-t) - \sum_{k=1}^4 \hat{\sigma}_{2k}^2 W_k^{\mathbb{Q}_1}(T-t) \right) \right\}, \quad (4.4.5)$$

where $\hat{\sigma} = \{\hat{\sigma}_{ij}\}$ is a matrix of constants so that the covariance matrix is given by $\hat{\sigma}\hat{\sigma}'$. By Girsanov's theorem the process $W^{\mathbb{Q}_2} = (W_1^{\mathbb{Q}_2}, \dots, W_4^{\mathbb{Q}_2})$ defined by

$$W_k^{\mathbb{Q}_2}(t) = W_k^{\mathbb{Q}_1}(t) - \hat{\sigma}_{2k} t, \quad (4.4.6)$$

is a Brownian motion under \mathbb{Q}_2 . As for S_b , we can use Ito's lemma to find the process for S_a under \mathbb{Q}_1 . In terms of $\hat{\sigma}$ we have

$$dS_a = S_a \left(\sum_{k=1}^4 \hat{\sigma}_{2k} (\hat{\sigma}_{2k} - \hat{\sigma}_{1k}) dt + \sum_{k=1}^4 (\hat{\sigma}_{1k} - \hat{\sigma}_{2k}) dW_k^{\mathbb{Q}_1} \right) \quad (4.4.7)$$

Using (4.4.6) in (4.4.7) we get the process for S_a under \mathbb{Q}_2 :

$$\begin{aligned} dS_a &= S_a \sum_{k=1}^4 (\hat{\sigma}_{1k} - \hat{\sigma}_{2k}) dW_k^{\mathbb{Q}_2} \\ &= S_a \sigma_a dW^{\mathbb{Q}_2}. \end{aligned} \quad (4.4.8)$$

Here $W^{\mathbb{Q}_2}$ is standard one-dimensional Brownian motion, and

$$\sigma_a = \sqrt{\sum_{k=1}^4 (\hat{\sigma}_{1k} - \hat{\sigma}_{2k})^2} = (\sigma_1^2 - 2\rho_{12}\sigma_1\sigma_2 + \sigma_2^2)^{\frac{1}{2}}.$$

We use (4.4.6) similarly to find the parameters of the process for S_b under \mathbb{Q}_2 . We find that

$$\begin{aligned} dS_b &= S_b \left(\sum_{k=1}^4 [\hat{\sigma}_{2k}(\hat{\sigma}_{3k} - \hat{\sigma}_{4k}) + \hat{\sigma}_{4k}(\hat{\sigma}_{4k} - \hat{\sigma}_{3k})] dt + \sum_{k=1}^4 (\hat{\sigma}_{3k} - \hat{\sigma}_{4k}) dW_k^{\mathbb{Q}_2} \right) \\ &= S_b (\mu_b dt + \sigma_b dW^{\mathbb{Q}_2}), \end{aligned} \quad (4.4.9)$$

where

$$\mu_b = \sigma_4^2 - \sigma_{34} + \sigma_{23} - \sigma_{24}$$

$$\text{and } \sigma_b = \bar{\sigma}_b.$$

Let $x = \ln(S_a(T)/S_a(t))$, $y = \ln(S_b(T)/S_b(t))$ and $\tau = T - t$. From (4.4.8) and (4.4.9) we see that the marginal densities under \mathbb{Q}_2 are given by $x \sim \mathcal{N}(-\frac{1}{2}\sigma_a^2\tau; \sigma_a\sqrt{\tau})$ and $y \sim \mathcal{N}((\mu_b - \frac{1}{2}\sigma_b^2)\tau; \sigma_b\sqrt{\tau})$. Let u and v be the standardized variables corresponding to x and y . The covariance between x and y is denoted $\sigma_{ab} = \rho_{ab}\sigma_a\sigma_b$. Then the conditional density for u and v corresponding to the condition $S_b > 1$ is

$$f(u|v > -d) = \int_{-d}^{\infty} f(u, v) dv = f(u) N\left(\frac{d + \rho_{ab}u}{\sqrt{1 - \rho_{ab}^2}}\right) \quad (4.4.10)$$

where $d = \frac{\ln S_b(t) + \mu_y}{\sigma_y}$ and $N(\cdot)$ denotes the cumulative standard normal distribution. Then the expectation in (4.4.4) is computed as

$$P(\mathbf{S}, t) = S_2(t) \int_{-\infty}^{\infty} (S_a(t) e^{u\sigma_x + \mu_x} - 1) f(u) N\left(\frac{d + \rho_{ab}u}{\sqrt{1 - \rho_{ab}^2}}\right) du \quad (4.4.11)$$

After the usual algebra we arrive at the following result:

Proposition 4. *Assume a generalized Black and Scholes economy with four assets following a four-dimensional geometric Brownian motion. Then the time t value of a European contingent claim promising the payoff*

$$P(\mathbf{S}, T) = (S_1 - S_2)\mathcal{X}_{(S_3 > S_4)},$$

at time T , is given by

$$P(\mathbf{S}, t) = S_1(t)N(d_1) - S_2(t)N(d_2), \quad (4.4.12)$$

where

$$\begin{aligned} d_1 &= \frac{\ln \frac{S_3(t)}{S_4(t)} - \left(\frac{1}{2}(\sigma_3^2 - \sigma_4^2) - \sigma_{13} + \sigma_{14}\right) \tau}{\sigma_b \sqrt{\tau}} \\ d_2 &= \frac{\ln \frac{S_3(t)}{S_4(t)} - \left(\frac{1}{2}(\sigma_3^2 - \sigma_4^2) - \sigma_{23} + \sigma_{24}\right) \tau}{\sigma_b \sqrt{\tau}} = d_1 - \rho_{ab} \sigma_a \sqrt{\tau} \\ \sigma_a &= (\sigma_1^2 - 2\rho_{12}\sigma_1\sigma_2 + \sigma_2^2)^{\frac{1}{2}} \\ \sigma_b &= (\sigma_3^2 - 2\rho_{34}\sigma_3\sigma_4 + \sigma_4^2)^{\frac{1}{2}} \\ \sigma_{ab} &= \sigma_{13} - \sigma_{14} + \sigma_{24} - \sigma_{23} \end{aligned}$$

Proof. The result follows from computing the expectation in (4.4.11). \square

Note that the formula in Proposition 4 degenerates to (4.3.2) if we let $S_3 = S_1$ and $S_4 = S_2$. Note also that the volatilities of the payout assets S_1 and S_2 do not enter into the pricing relation. Only their covariance with the trigger assets matter.

The following corollary will be useful:

Corollary 1. *Let g and h be positive constants, and let $P(gS_1, hS_2, S_3, S_4; t)$ denote the current value of a four-asset exchange option. Then*

$$P(gS_1, hS_2, S_3, S_4; t) = gS_1N(d_1) - hS_2N(d_2), \quad (4.4.13)$$

where d_1 and d_2 are given as in Proposition 1.

Proof. Follows from the additive structure of (4.4.12), and because d_1 and d_2 are independent of S_1 and S_2 . \square

For the sake of completeness, we conclude this subsection by showing the value of a contingent exchange option, i.e., the holder receives the non-negative part of the difference between S_1 and S_2 if S_3 is larger than S_4 .

Proposition 5. *The time t value of a European derivative promising the payoff*

$$P(\mathbf{S}, T) = (S_1 - S_2)^+ \mathcal{X}_{(S_3 > S_4)},$$

at time T is given by

$$P(\mathbf{S}, t) = S_1(t)N_2(d_1, \hat{d}_1, \rho_{ab}) - S_2(t)N_2(d_2, \hat{d}_2, \rho_{ab}),$$

where $N_2(\cdot)$ denotes the bivariate cumulative standard normal distribution and

$$\begin{aligned} d_1 &= \frac{\ln \frac{S_1(t)}{S_2(t)} + \frac{1}{2}\sigma_a\tau}{\sigma_a\sqrt{\tau}} \\ \hat{d}_1 &= \frac{\ln \frac{S_3(t)}{S_4(t)} - \left(\frac{1}{2}(\sigma_3^2 - \sigma_4^2) - \sigma_{13} + \sigma_{14}\right)\tau}{\sigma_b\sqrt{\tau}} \\ d_2 &= \frac{\ln \frac{S_1(t)}{S_2(t)} - \frac{1}{2}\sigma_a\tau}{\sigma_a\sqrt{\tau}} = d_1 - \sigma_a\sqrt{\tau} \\ \hat{d}_2 &= \frac{\ln \frac{S_3(t)}{S_4(t)} - \left(\frac{1}{2}(\sigma_3^2 - \sigma_4^2) - \sigma_{23} + \sigma_{24}\right)\tau}{\sigma_b\sqrt{\tau}} = \hat{d}_1 - \rho_{ab}\sigma_a\sqrt{\tau} \\ \sigma_a &= (\sigma_1^2 - 2\rho_{12}\sigma_1\sigma_2 + \sigma_2^2)^{\frac{1}{2}} \\ \sigma_b &= (\sigma_3^2 - 2\rho_{34}\sigma_3\sigma_4 + \sigma_4^2)^{\frac{1}{2}} \\ \sigma_{ab} &= \sigma_{13} - \sigma_{14} + \sigma_{24} - \sigma_{23} \end{aligned}$$

Proof. We have

$$\begin{aligned} & \mathbb{E}^{\mathbb{Q}_1} \left[e^{-r(T-t)} (S_1 - S_2)^+ \mathcal{X}_{(S_3 > S_4)} \right] \\ &= S_2(t) \int_{-\infty}^{\infty} (S_a(t)e^{u\sigma_x + \mu_x} - 1)^+ f(u) N \left(\frac{d + \rho_{ab}u}{\sqrt{1 - \rho_{ab}^2}} \right) du \\ &= S_2(t) \int_{-d^*}^{\infty} (S_a(t)e^{u\sigma_x + \mu_x} - 1) f(u) N \left(\frac{d + \rho_{ab}u}{\sqrt{1 - \rho_{ab}^2}} \right) du, \end{aligned} \tag{4.4.14}$$

where $d^* = (\ln S_b + \mu_x) \frac{1}{\sigma_x}$, by the same arguments that take us from (4.4.2) to (4.4.11). After a change of variable from u to $k = u - \sigma_x$, equation (4.4.14) can be written as

$$S_2(t) \left[S_a e^{\mu_x + \frac{\sigma_x^2}{2}} \int_{-\infty}^{\hat{d}^*} f(k) N \left(\frac{\hat{d} + \rho_{ab} k}{\sqrt{1 - \rho_{ab}^2}} \right) dk - \int_{-\infty}^{d^*} f(u) N \left(\frac{d + \rho_{ab} u}{\sqrt{1 - \rho_{ab}^2}} \right) du \right] \quad (4.4.15)$$

where $\hat{d}^* = (\ln S_a + \mu_x + \sigma_x^2) \frac{1}{\sigma_x}$ and $\hat{d} = d + \rho_{ab} \sigma_x$. Substituting variables and computing the expectations in (4.4.15) yields the desired result. □

4.4.2 Assuming a Joint Lognormal Distribution

We are now ready to return to the decomposition of the portfolio exchange option's payoff in (4.4.1). Armed with Proposition 1 and Corollary 1, we can obtain an expression for the current value of (4.4.1). We substitute the variables S_i in Proposition 1 in the following way:

$$\begin{aligned} S_1 &\equiv S_i & S_2 &\equiv S_j \\ S_3 &\equiv I_1 = \sum_{i=1}^k S_i & S_4 &\equiv I_2 = \sum_{j=k+1}^n S_j \end{aligned}$$

Proposition 6. *Suppose assets are joint lognormally distributed. Pretend the portfolio values $I_1(T) = \sum_{i=1}^k S_i(T)$ and $I_2(T) = \sum_{j=k+1}^n S_j(T)$ are lognormal as well. Then the value of a derivative of European type promising the payoff*

$$P(\mathbf{S}, T) = (I_1(T) - I_2(T))^+,$$

at time T is given by

$$P(\mathbf{S}, t) = \sum_{i=1}^k S_i(t) N(d_{1i}) - \sum_{j=k+1}^n S_j(t) N(d_{2j}), \quad (4.4.16)$$

where $N(\cdot)$ denotes the cumulative standard normal distribution and

$$\begin{aligned}
 d_{1i} &= \frac{\ln \frac{I_1(t)}{I_2(t)} - \left(\frac{1}{2}(\sigma_{I_1}^2 - \sigma_{I_2}^2) - \sigma_{iI_1} + \sigma_{iI_2}\right) \tau}{\sigma_b \sqrt{\tau}} \\
 d_{2j} &= \frac{\ln \frac{I_1(t)}{I_2(t)} - \left(\frac{1}{2}(\sigma_{I_1}^2 - \sigma_{I_2}^2) - \sigma_{jI_1} + \sigma_{jI_2}\right) \tau}{\sigma_b \sqrt{\tau}} = d_{1i} - \rho_{ab} \sigma_a \sqrt{\tau} \\
 \sigma_a &= (\sigma_i^2 - 2\rho_{ij} \sigma_i \sigma_j + \sigma_j^2)^{\frac{1}{2}} \\
 \sigma_b &= (\sigma_{I_1}^2 - 2\rho_{I_1 I_2} \sigma_{I_1} \sigma_{I_2} + \sigma_{I_2}^2)^{\frac{1}{2}} \\
 \sigma_{ab} &= \sigma_{iI_1} - \sigma_{iI_2} + \sigma_{jI_2} - \sigma_{jI_1}
 \end{aligned}$$

Proof. From Proposition 4 and Corollary 1 we know the value of each element of the double sum in (4.4.1). Hence,

$$P(\mathbf{S}, t) = \sum_{i=1}^k \sum_{j=k+1}^n \left[\frac{1}{n-k} S_i N(d_{1i}) - \frac{1}{k} S_j N(d_{2j}) \right] \quad (4.4.17)$$

because the price of a sum of claims is the sum of the prices of each claim. The double sum in (4.4.17) simplifies to (4.4.16). □

The parameters σ_I and σ_{iI} must be approximated. Let $I = \sum_{j=1}^s S_j$. Noting that

$$\begin{aligned}
 \sigma_{iI} = \text{cov}(\tau_i, \tau_I) &= \text{cov} \left(\ln \frac{S_i(T)}{S_i(0)}, \ln \frac{I(T)}{I(0)} \right) \\
 &= \text{cov}(\ln S_i, \ln I) \\
 &= \text{cov}(\ln S_i, \ln(S_1 + S_2 + \dots)),
 \end{aligned}$$

we see the need for approximating the logarithm of a sum. A standard approach is to approximate the real portfolio return $\tau_I = \ln \sum_{i=1}^s w_i e^{r_i}$ with $\tau_I \approx \sum_{i=1}^s w_i r_i$ where $w_i = \frac{S_i}{\sum_{j=1}^s S_j}$. Consequently we have

$$\sigma_{iI} \approx \sum_{j=1}^s w_j \sigma_{ij}.$$

Moreover, the volatility of the portfolio σ_I is correspondingly approximated as

$$\sigma_I \approx \left(\sum_{i=1}^s \sum_{j=1}^s w_i w_j \sigma_{ij} \right)^{\frac{1}{2}}.$$

Note that we have used the same assumptions regarding the distribution of I_1 and I_2 in our case as in the standard approach of (4.3.2). The parameters of the portfolio process are

Table 4.3: Test results for Proposition 6.

	Monte Carlo	Stand. (4.3.2)	Prop. 6 (4.4.16)
Price	12.8515	12.6025	12.5998

also essentially approximated in the same way. Yet the price obtained by using (4.4.16) differs from that of (4.3.2), as can be seen from Table 4.3. Prices are calculated on the basis of the assumptions in Table 4.1. The reason for this is that in our approach we are able to use information about each asset's distributional properties, in particular the covariance σ_{iI} , rather than only relying on the distributional properties of the portfolios. The difference in Table 4.3 is not large. In section 4.6 we test for the approach that produces the best general approximation, for various parameters.

The main advantage of the current approach, however, is the possibility of letting the portfolio have a different distribution from that of the individual assets. An example is given in the next section.

4.5 Pricing Under a Joint Lognormal-Normal Distribution

It would have been possible to develop a formula similar to (4.3.2), the standard approximation, also if we had assumed that I_1 and I_2 were normally distributed rather than lognormal as in (4.3.2). The distribution of the individual assets would not then have been needed.

However, if nothing is said about the prices of individual assets, a risk-neutral valuation relationship has to be deduced on the basis of the preferences of a representative agent and the distribution of the market portfolio, (see Rubinstein (1976) and Brennan (1979)). In our setting, where active managers have subjective views on individual assets and hold the difference between two portfolios actively, one could argue that such an approach would be non-consistent. Hence, a model with well-defined properties for individual prices is more desirable.

By separating payout assets and trigger assets by means of the decomposition in (4.4.1), we now develop a pricing formula in which the generalized Black and Scholes model applies to the payout assets, i.e., they are joint lognormally distributed. The trigger asset however, which is the difference between the two portfolios, is assumed to be normally distributed. The joint distribution of assets and difference between portfolios is a joint lognormal-normal distribution.

The true distributions of I_1 and I_2 are unknown in a generalized Black and Scholes model. The central limit theorem can not be used because the variables I_i are always non-negative, being sums of S_i^t s, which is non-negative by assumption. And, more importantly, the prices of assets are generally not independent.

However, the difference between I_1 and I_2 may be negative. Moreover, the difference may be seen as the sum of long-short positions each with an expectation that is positive but close to zero, and different volatilities. According to Grinold and Kahn (1995), who claims that to maximize his risk-adjusted active return the manager seeks to find many independent bets, the correlation between the long-short positions might be low. It may thus be more appealing to assume a normal distribution for this difference than for each portfolio separately, even though the independence criteria still prohibit an outright application of the central limit theorem.

Consider the following reformulation of the condition, or trigger, in (4.4.1):

$$\begin{aligned} I_1(T) &> I_2(T) \\ I_1(t) \frac{I_1(T)}{I_1(t)} &> I_2(t) \frac{I_2(T)}{I_2(t)} \\ q \frac{I_1(T)}{I_1(t)} - \frac{I_2(T)}{I_2(t)} &> 0 \end{aligned} \tag{4.5.1}$$

$$\eta > 1 - q \tag{4.5.2}$$

where $\eta = qr_1 - r_2$, q is the constant $q = \frac{I_1(t)}{I_2(t)}$, and r_i is the arithmetic return of I_i . In this section we assume that $I_1(T) - I_2(T)$, and consequently η , is normally distributed under the risk-adjusted probability measure \mathbb{Q}_1 . Hence, η , the difference between the arithmetic returns, can be described by a Brownian motion as follows:

$$\eta = r(q - 1)dt + \sigma_\eta dW^{\mathbb{Q}_1}, \tag{4.5.3}$$

where $\sigma_\eta = (q^2\sigma_{r_1}^2 + \sigma_{r_2}^2 - 2q\sigma_{r_1r_2})^{\frac{1}{2}}$ by assumption. σ_{r_i} is the standard deviation of the arithmetic portfolio returns. Hence, the marginal distribution of η under the risk-adjusted measure is $\eta \sim \mathcal{N}(r(q - 1)\tau, \sigma_\eta\sqrt{\tau})$.

Consequently, the assets S_i and the difference between the two portfolios are multivariate lognormally-normally distributed. In particular, the two-dimensional variate $\left(\frac{S_a(T)}{S_a(t)}, \eta\right)$, where as before $S_a = \frac{S_i}{S_j}$, is bivariate lognormally-normally distributed. Note that (4.5.3) holds purely by assumption. It can be considered as a new primitive of the model. Because we do not know the process for I_i under \mathbb{P} we can not deduce the process under \mathbb{Q}_1 either, it must be assumed, as exemplified by (4.5.3).

The joint lognormal-normal, or semi-lognormal, density function is described by Crow and Shimizu (1988) and has been used earlier by Camara and Stapleton (1998) where the pricing kernel is studied under the assumption of lognormally-normally distributed assets and aggregate wealth.

Let u and v be standardized variables, so that $u = \frac{x - \mu_x}{\sigma_x}$ and $v = \frac{\eta - \mu_\eta}{\sigma_\eta}$. The variable x has been defined earlier as $\ln \frac{S_a(T)}{S_a(t)}$. Hence $(u, v) = \left(\ln \frac{S_a(T)}{S_a(t)}, \eta \right)_{stand.}$ follows the joint standard normal distribution with density function

$$f(u, v) = \frac{1}{2\pi\sqrt{1 - \rho_{vu}^2}} e^{-\frac{1}{2(1 - \rho_{vu}^2)}[u^2 - 2\rho_{vu}uv + v^2]}. \quad (4.5.4)$$

The assumptions in this section do not affect the form of the pricing kernel, i.e., $\frac{dQ_1}{dP}$, nor is our earlier definition of $\frac{dQ_2}{dQ_1}$ affected. Hence, Girsanov's theorem can be used as before to find the distribution for x and η under Q_2 .

Proposition 7. *Suppose assets are joint lognormally distributed. Suppose assets and differences between two portfolios composed of these assets, i.e., long-short portfolios, follow a joint lognormal-normal distribution. In particular, assume that the stochastic properties of η under the risk-adjusted measure are described by (4.5.3). Then the value of a derivative of European type promising the payoff*

$$P(\mathbf{S}, T) = (I_1(T) - I_2(T))^+,$$

at time T , where $I_1(T) = \sum_{i=1}^k S_i(T)$ and $I_2(T) = \sum_{j=k+1}^n S_j(T)$, is given by

$$P(\mathbf{S}, t) = \sum_{i=1}^k S_i(t)N(d_{1i}^*) - \sum_{j=k+1}^n S_j(t)N(d_{2j}^*), \quad (4.5.5)$$

where $N(\cdot)$ denotes the cumulative standard normal distribution and

$$\begin{aligned} d_{1i}^* &= \frac{\left(\frac{I_1(t)}{I_2(t)} - 1\right)(1 + r\tau) + \sigma_{\eta i}\tau}{\sigma_\eta\sqrt{\tau}} \\ d_{2j}^* &= \frac{\left(\frac{I_1(t)}{I_2(t)} - 1\right)(1 + r\tau) + \sigma_{\eta j}\tau}{\sigma_\eta\sqrt{\tau}} = d_{1i}^* - \rho_{\eta x}\sigma_x\sqrt{\tau} \\ \sigma_x &= (\sigma_i^2 - 2\sigma_{ij} + \sigma_j^2)^{\frac{1}{2}} \\ \sigma_{\eta i} &= \frac{I_1(t)}{I_2(t)}\sigma_{ir_1} - \sigma_{ir_2} \\ \sigma_\eta &= (q^2\sigma_{r_1}^2 + \sigma_{r_2}^2 - 2q\sigma_{r_1r_2})^{\frac{1}{2}} \end{aligned}$$

Proof. Let

$$\begin{aligned} P\left(\frac{1}{n-k}S_i, \frac{1}{k}S_j, \eta; t\right) &= \mathbb{E}^{\mathbb{Q}_1} \left[e^{-r\tau} \left(\frac{1}{n-k}S_i(T) - \frac{1}{k}S_j(T) \right) \mathcal{X}_{(\eta > 1-q)} \right] \\ &= \frac{1}{k}S_j(t) \mathbb{E}^{\mathbb{Q}_2} \left[\left(\frac{k}{n-k}S_a - 1 \right) \mathcal{X}_{(\eta > 1-q)} \right], \end{aligned}$$

where η is given in (4.5.2). There exists a 3×3 matrix $\hat{\sigma}$ such that $\hat{\sigma}\hat{\sigma}$ is the covariance matrix between the returns of asset i , asset j and η . Then $\sigma_\eta dW^{\mathbb{Q}_1}$, where $dW^{\mathbb{Q}_1}$ is a correlated Wiener process, can be written as $\sum_{m \in \{i, j, \eta\}} \hat{\sigma}_{\eta m} dW_m^{\mathbb{Q}_1}$, where $dW_m^{\mathbb{Q}_1}$ is independent Wiener processes for $m \in \{i, j, \eta\}$.

Using Girsanov's theorem and $\frac{d\mathbb{Q}_2}{d\mathbb{Q}_1}$ given in (4.4.5) we have

$$\begin{aligned} \eta &= r(q-1)dt + \sigma_\eta dW^{\mathbb{Q}_1} \\ &= r(q-1)dt + \sum_{m \in \{i, j, \eta\}} \hat{\sigma}_{\eta m} dW_m^{\mathbb{Q}_1} \\ &= r(q-1)dt + \sum_{m \in \{i, j, \eta\}} \hat{\sigma}_{\eta m} \left(dW_m^{\mathbb{Q}_2} + \hat{\sigma}_{mj} dt \right) \end{aligned} \quad (4.5.6)$$

Noting that $\sum_{m \in \{i, j, \eta\}} \hat{\sigma}_{\eta m} \hat{\sigma}_{mj} = \sigma_{\eta j}$, that is, the covariance between η and the return of asset S_j , we see that under \mathbb{Q}_2 , $\mu_\eta = r(q-1) + \sigma_{\eta j}$. The standard deviation of η is not affected by the change of measure. Thus, under \mathbb{Q}_2 , the marginal distribution for η is normal: $\eta \sim \mathcal{N}(\mu_\eta \tau, \sigma_\eta \sqrt{\tau})$. The marginal distribution for x is normal with the same parameters as those in Proposition 1.

From (4.5.4) we see that the conditional distribution for u given v , corresponding to the condition $\eta > 1 - k$, is

$$f(u|v > -d^*) = \int_{-d^*}^{\infty} f(u, v) dv = f(u) \mathbb{N} \left(\frac{d^* + \rho_{\eta x} u}{\sqrt{1 - \rho_{\eta x}^2}} \right) \quad (4.5.7)$$

where $d^* = \frac{q-1+\mu_\eta \tau}{\sigma_\eta \sqrt{\tau}}$. Note that $\rho_{vu} = \rho_{\eta x}$. Then

$$\begin{aligned} &\mathbb{E}^{\mathbb{Q}_1} \left[e^{-r\tau} P\left(\frac{1}{n-k}S_i, \frac{1}{k}S_j, \eta; T\right) \right] \\ &= \frac{1}{k}S_j(t) \int_{-\infty}^{\infty} \left(\frac{k}{n-k}S_a(t) e^{u\sigma_x + \mu_x} - 1 \right) f(u) \mathbb{N} \left(\frac{d^* + \rho_{\eta x} u}{\sqrt{1 - \rho_{\eta x}^2}} \right) du \\ &= \frac{1}{n-k}S_i(t) \mathbb{N}(d_1^*) - \frac{1}{k}S_j(t) \mathbb{N}(d_2^*), \end{aligned} \quad (4.5.8)$$

Table 4.4: Test results of Proposition 7. Base Case.

	Monte Carlo	Stand.	Prop. 3	Prop. 7
Price	12.8515	12.6025	12.5998	12.5716

where $d_1^* = \left[\left(\frac{I_1(t)}{I_2(t)} - 1 \right) (1 + r\tau) + \sigma_{\eta_i}\tau \right] \frac{1}{\sigma_{\eta}\sqrt{\tau}}$ and $d_2^* = d_1^* - \rho_{\eta x}\sigma_x\sqrt{\tau}$.

Note that d_1^* does not depend on asset j or the parameters of its process. Likewise d_2^* does not depend on asset i in any way. Hence,

$$\sum_{i=1}^k \sum_{j=k+1}^n P\left(\frac{1}{n-k}S_i, \frac{1}{k}S_j, \eta; t\right)$$

simplifies to (4.5.5).

□

The calculation of σ_{r_i} and σ_{jr_1} for $i \in \{1, 2\}$, requires comment. We have $\sigma_{jr_1} = \text{cov}(\ln S_j, r_1)$. Because

$$r_1 = \frac{I_1(T)}{I_1(t)} - 1 = \sum_{i=1}^k w_i(1 + r_{S_i})$$

where $w_i = \frac{S_i}{\sum_{i=1}^k S_i}$ and $r_{S_i} = \frac{S_i(T)}{S_i(t)} - 1$, we have

$$\sigma_{jr_1} = \sum_{i=1}^k w_i \text{cov}\left(\ln \frac{S_j(T)}{S_j(t)}, \frac{S_i(T)}{S_i(t)}\right).$$

Hence σ_{jr_1} is a sum of covariances between logarithmic return for asset j and arithmetic return for asset i . To be consistent with the assumption in (4.5.3), the volatility for the portfolio, σ_{r_1} , is computed using a covariance matrix of arithmetic returns.

However, in Table 4.4 the value of (4.5.5) is computed on the basis of the parameters of Table 4.1, even though that covariance matrix is based on logarithmic returns. This is done to isolate the effects of computational differences in the pricing formulas. Moreover, the difference between logarithmic and arithmetic covariance is very small. The prices in Table 4.4 indicate that there is little difference between the standard approach and the two pricing formulas presented here. Comparisons over a broader range of parameters will be presented in section 4.6.

The formula in Proposition 7 becomes quite simple when the option is "at the money", i.e., when $I_1(t) = I_2(t)$ or equivalently when $q = 1$. In that case σ_{η} is the volatility of the

difference between $I_1(t)$ and $I_2(t)$. It may be interpreted as the gross volatility of an active portfolio.¹

Corollary 2. *Suppose $I_1(t) = I_2(t)$ in Proposition 7. Then the current value of the payoff $(I_1(T) - I_2(T))^+$ is*

$$P(\mathbf{S}; t) = \sum_{i=1}^k S_i(t)N(\beta_{i\eta}\sigma_\eta\sqrt{\tau}) - \sum_{j=k+1}^n S_j(t)N(\beta_{j\eta}\sigma_\eta\sqrt{\tau}) \quad (4.5.9)$$

where $\beta_{i\eta} = \frac{\sigma_{i\eta}}{\sigma_\eta^2}$ is the beta of asset i with respect to the active portfolio, i.e., holdings of $I_1(t)$ less holdings of $I_2(t)$.

Proof. Follows from Proposition 7, assuming $q = 1$. □

4.5.1 Pricing a Portfolio Spread Option

An extension of Proposition 7 to the case of a portfolio spread option is readily available. A spread option is an exchange option with non-zero strike, i.e., we consider a contingent claim with payoff

$$\max(I_1(T) - I_2(T), K), \quad (4.5.10)$$

where K is the strike. We will allow $K \in \mathbb{R}$. Normally an active manager will accept some potential downside and choose $K < 0$.

Proposition 8. *Adopt the assumptions of Proposition 7. Consider a derivative with payoff $\max(I_1(T) - I_2(T), K)$. Its value at time t is*

$$P(\mathbf{S}, t) = \sum_{i=1}^k S_i(t)N(d_{1i}^*) - \sum_{j=k+1}^n S_j(t)N(d_{2j}^*) + e^{-r\tau}KN(d_K^*), \quad (4.5.11)$$

¹Note that this is different from *relative volatility* or *tracking error*, which is the volatility of active returns relative to a benchmark. σ_η is calculated on the basis of active positions only, independently of any benchmarkholdings.

where $N(\cdot)$ denotes the cumulative standard normal distribution and

$$\begin{aligned}
 d_{1i}^* &= \frac{\left(\frac{I_1(t)}{I_2(t)} - 1\right) (1 + r\tau) - \frac{K}{I_2(t)} + \sigma_{\eta i} \tau}{\sigma_{\eta} \sqrt{\tau}} \\
 d_{2j}^* &= \frac{\left(\frac{I_1(t)}{I_2(t)} - 1\right) (1 + r\tau) - \frac{K}{I_2(t)} + \sigma_{\eta j} \tau}{\sigma_{\eta} \sqrt{\tau}} = d_{1i}^* - \rho_{\eta x} \sigma_x \sqrt{\tau} \\
 d_K^* &= \frac{\frac{K}{I_2(t)} - \left(\frac{I_1(t)}{I_2(t)} - 1\right) (1 + r\tau)}{\sigma_{\eta} \sqrt{\tau}}
 \end{aligned} \tag{4.5.12}$$

All parameters defined as in Proposition 7.

Proof. First note that the value of the payoff in (4.5.10) can be formulated in the following way:

$$\mathbb{E}^{\mathbb{Q}_1} \left[e^{-r\tau} \left\{ (I_1(T) - I_2(T)) \mathcal{X}_{(\eta > 1 - q + \frac{K}{I_2(t)})} + K \mathcal{X}_{(\eta \leq 1 - q + \frac{K}{I_2(t)})} \right\} \right] \tag{4.5.13}$$

This expectation of a sum can be calculated as two separate expectations. The value of the first part of the sum in (4.5.13) follows immediately from the proof of Proposition 7, when the condition (4.5.2) is replaced with $\eta > 1 - q + \frac{K}{I_2(t)}$, which is equivalent to the inequality $I_1(T) - I_2(T) > K$. The second element of the sum in (4.5.13) is easily evaluated when the fact that $\mathbb{E}^{\mathbb{Q}_1} \left[\mathcal{X}_{(\eta \leq 1 - q + \frac{K}{I_2(t)})} \right]$ is equal to $\mathbb{Q}_1 \left(\eta \leq 1 - q + \frac{K}{I_2(t)} \right)$ is recognized. \square

A closed form solution for the value of the payoff (4.5.10) is not available when the portfolio values are assumed to be lognormally distributed. In Zhang (1997) a survey of different numerical solutions is given. Numerical testing of Proposition 8 is provided in the next section.

4.6 Numerical Testing

In this section we numerically test the various approaches discussed earlier. We compare prices obtained with the standard approach in (4.3.2) with the results obtained using Proposition 6, Proposition 7 and Proposition 8. All values are compared with the true Monte Carlo price of the option as well as the standard approximation.

Because of the potentially huge number of perturbations of parameters in these formulae, an exhaustive test of the goodness of each formula is not possible. However, to get a feel for

Table 4.5: Approximated prices, at-the-money.

	Monte Carlo	Standard	Prop. 6	Prop. 7
Case 1	12.8515	12.6026	12.5998	12.5702
Case 2	15.5600	15.2464	15.2425	15.1727
Case 3	13.7015	12.6025	12.5998	12.5716
Case 4	10.4078	10.2032	10.2015	10.1942
Case 5	12.7547	12.5079	12.5047	12.4557
Case 6	14.9541	14.6375	14.6330	14.6153
Case 7	11.0426	10.8441	10.8355	10.7623
Average Bias	-	-1.98%	-2.01%	-2.29%
Std.dev. Bias	-	0.14%	0.12%	0.19%

how well the pricing formula performs with various combinations of parameters, we proceed by choosing a limited number of cases.

The first case is chosen as the base case given by Table 4.1 earlier. Then the volatilities of the assets in I_1 are increased in case 2 and the volatilities of the assets in I_2 are increased in case 3. In cases 4 and 5 we reduce the volatilities of I_2 and I_1 respectively. In the sixth case, all correlations are lowered to zero, and in case 7 all correlations are increased. The parameters of the different cases are given in Appendix A.

In all the seven cases there are two assets in both I_1 and I_2 , and they all have the same initial value. Hence, at the present we are studying "at-the-money" options. The maturity of the option is one year.

Table 4.5 compares option prices computed using the different formulae in the seven different cases. The bias is computed as the average over a percentage difference from the Monte Carlo value.

The value obtained by using the formula in Proposition 6 is included for the sake of completeness. Not surprisingly, there is very little difference between this value and the standard formula, as they are based on the same assumptions.

The candidate for a better pricing approximation is the formula in Proposition 7, which is based on the joint lognormal-normal, or semi-lognormal, assumption. However, Table 4.5 indicates that this is not a more realistic assumption. The pricing formula in Proposition 7 performs slightly more poorly than the more simple standard approximation of equation (4.3.2), when at-the-money options are valued.

A striking observation from Table 4.5 is the consistency of the negative bias in all pricing formulae. It does not vary much in size across different sets of parameters. This seems to be the case also when the prices of individual assets vary, but when the condition $I_1(t) = I_2(t)$ is still kept. However, when that condition does not hold, the bias starts to vary. Typically, Proposition 4 overvalues the option when $I_2(t) > I_1(t)$, while the standard approximation keeps undervaluing albeit to a variable degree. When $I_1(t) < I_2(t)$ both approximations undervalue the option to a variable degree, depending on *how* $I_1(t)$ is less than $I_2(t)$, i.e., what composition of $I_1(t)$ makes it less than $I_2(t)$. These calculations are carried out using the parameters of Case 1, but the initial values of the assets has been changed in various ways. An active manager, however, will usually compose his active portfolio so that under-weights are matched with over-weights, i.e., so that $I_1(t) = I_2(t)$.

The assumption of semi-lognormality is perhaps more plausible when the portfolios I_1 and I_2 contain many assets. To investigate this hypothesis, we simulated a varying number of assets for a given set of parameters. In Table 4.6, we vary the total number of assets from 4 to 100, always using half the number of assets in each of the portfolios I_1 and I_2 . The initial values of each asset are the same for all assets and are reduced as the number of assets is increased. Thus, the value of the portfolio is the same in all cases, and the initial portfolio values are 200. The calculations are based on a correlation coefficient of 0.2 between all assets, and volatilities varying from 4% to 25%, evenly spread across the N assets. The number of assets is given in the left column. From Table 4.6 we see a slight improvement in pricing accuracy for both the standard approximation and Proposition 7 when the number of assets in I_1 and I_2 increases. The difference between the two approximations is reduced with an increasing number of assets.

The reduction in option price as the number of assets increases is mainly due to a diversification effect. When the number of assets is increased from 4 to 100, σ_η decreases from 17.03% to 2.78%. The notional amounts invested in I_1 and I_2 are the same in all cases.

Table 4.6: Differing number of assets - N.

N	Monte Carlo	Standard	Prop. 7	Bias Stand.	Bias Prop. 7
4	13.8574	13.5711	13.5403	-2.07%	-2.29%
8	8.2002	8.1515	8.1408	-0.59%	-0.72%
16	5.6021	5.6806	5.6787	1.40%	1.37%
32	3.9100	3.8922	3.8914	-0.46%	-0.48%
64	2.8454	2.8156	2.8153	-1.05%	-1.06%
100	2.2301	2.2181	2.2180	-0.54%	-0.54%

Table 4.7: Approximated prices, varying strike - K.

K	Monte Carlo	Prop. 5	Bias
25	27.5960	27.3954	-0.73%
20	24.0981	23.8704	-0.94%
15	20.8647	20.6126	-1.21%
10	17.9066	17.6385	-1.50%
5	15.2372	14.9573	-1.84%
0	12.8515	12.5702	-2.19%
-5	10.7615	10.4704	-2.71%
-10	8.9395	8.6450	-3.29%
-15	7.3652	7.0753	-3.94%
-20	6.0275	5.7394	-4.78%
-25	4.8983	4.6133	-5.82%

Table 4.7 shows the pricing bias of the formula for the *portfolio spread option* in Proposition 8. We have assumed $I_1 = I_2$, and the parameters of Case 1 has been adopted. The strike varies in value as shown in the left column of the table. Here the pricing bias is not consistent. From Table 4.7 we see the bias reduced when K is positive and increasing. However, Proposition 8 seems to price less accurately when K is increasingly negative.

4.7 Concluding Remarks

In this chapter we have developed a pricing formula for portfolio spread options when assets and differences between portfolios of assets are joint lognormally-normally distributed. Numerical testing shows that the pricing formula does not produce a more accurate value than a standard approximation, when the true process for the assets is multi-dimensional lognormal. For at-the-money options with zero strike, the numerical tests show a consistent negative bias for

various sets of parameters and a given number of assets.

The separation of payout asset and trigger assets used in this chapter provides a method to develop other pricing formulas on the basis of joint distributions with differing properties along certain dimensions. The joint lognormal-normal distribution is only one of several examples of such distributions, as shown in Crow and Shimizu (1988).

In Propositions 7 and 8 the pricing formulas are correct only to the extent that the distributional assumptions for assets and portfolios are consistent with reality. An alternative approach would be to approximate the payoff, which could then be priced correctly because only the distribution of individual assets would have been involved. Hence, rather than the approach followed in Propositions 7 and 8 where approximating formulas for the right payoff are developed, we could develop a correct formula for an approximated payoff. This approach we leave for further studies.

4.8 Appendix: Parameters of the Seven Cases

Table 4.8: Case 1 - Initial case

Asset	S_1	S_2	S_3	S_4
$S_i(0)$	100	100	100	100
S_1	0.10			
S_2	0.2	0.15		
S_3	0.2	0.3	0.20	
S_4	0.1	0.4	0.2	0.25

Table 4.9: Case 2 - High volatility of I_2

Asset	S_1	S_2	S_3	S_4
$S_i(0)$	100	100	100	100
S_1	0.10			
S_2	0.2	0.15		
S_3	0.2	0.3	0.25	
S_4	0.1	0.4	0.2	0.30

Table 4.10: Case 3 - High volatility of I_1

Asset	S_1	S_2	S_3	S_4
$S_i(0)$	100	100	100	100
S_1	0.15			
S_2	0.2	0.20		
S_3	0.2	0.3	0.20	
S_4	0.1	0.4	0.2	0.25

Table 4.11: Case 4 - Low volatility of I_2

Asset	S_1	S_2	S_3	S_4
$S_i(0)$	100	100	100	100
S_1	0.10			
S_2	0.2	0.15		
S_3	0.2	0.3	0.15	
S_4	0.1	0.4	0.2	0.20

Table 4.12: Case 5 - Low volatility of I_1

Asset	S_1	S_2	S_3	S_4
$S_i(0)$	100	100	100	100
S_1	0.05			
S_2	0.2	0.1		
S_3	0.2	0.3	0.20	
S_4	0.1	0.4	0.2	0.25

Table 4.13: Case 6 - Low correlation

Asset	S_1	S_2	S_3	S_4
$S_i(0)$	100	100	100	100
S_1	0.10			
S_2	0.0	0.15		
S_3	0.0	0.0	0.20	
S_4	0.0	0.0	0.0	0.25

Table 4.14: Case 7 - High correlation

Asset	S_1	S_2	S_3	S_4
$S_i(0)$	100	100	100	100
S_1	0.10			
S_2	0.5	0.15		
S_3	0.7	0.8	0.20	
S_4	0.6	0.4	0.6	0.25

Bibliography

- ADLER, M. AND J. DETEMPLE (1988): "On the Optimal Hedge of a Nontraded Cash Position," *Journal of Finance*, 43:143–153.
- BARDHAN, I. (1995): "Exchange Rate Shocks, Currency options and the Siegel Paradox," *Journal of International Money and Finance*, 14(3):441–458.
- BICK, A. (1987): "On the Consistency of the Black-Scholes Model with a General Equilibrium Framework," *Journal of Financial and Quantitative Analysis*, 22:259–275.
- (1990): "On Viable Diffusion Price Processes of the Market Portfolio," *Journal of Finance*, 2000:673–689.
- BLACK, F. (1989): "Universal Hedging: Optimizing Currency Risk and Reward in International Equity Portfolios," *Financial Analysts Journal*, (3000):16–22.
- (1990): "Equilibrium Exchange Rate Hedging," *Journal of Finance*, 45:899–907.
- BLACK, F. AND M. SCHOLES (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81(3):637–654.
- BORCH, K. (1962): "Equilibrium in a Reinsurance Market," *Econometrica*, 30:424–444.
- BOYER, R. S. (September 1972): "The Interest Rate Parity Theorem Again: A Resolution of the Siegel Paradox," mimeo, University of Western Ontario.
- BRENNAN, M. (1979): "The Pricing of Contingent Claims in Discrete Time Models," *The Journal of Finance*, 24:53–68.
- BRENNAN, M. AND R. SOLANKI (1981): "Optimal Portfolio Insurance," *Journal of Financial and Quantitative Analysis*, 16:279–300.

- BRIYS, E. AND S. BRUNO (1992): "Optimal Currency Hedge Ratios and Interest Rate Risk," *Journal of International Money and Finance*, 11:431-445.
- BRIYS, E., M. CROUHY, AND H. SCHLESINGER (1993): "Optimal Hedging in a Futures Market with Background Noise and Basis Risk," *European Economic Review*, 37:949-960.
- CAMARA, A. AND R. C. STAPLETON (1998): "Option Pricing Under Constant Absolute Risk Aversion and Constant Proportional Risk Aversion," Working Paper, University of Strathclyde, Department of Accounting and Finance.
- CARR, P. AND D. MADAN (1997): "Optimal Positioning in Derivative Securities," University of Maryland, College Park, MD 20742.
- COX, J. AND C. HUANG (1989): "Optimal consumption and Portfolio policies when Asset Prices Follow a Diffusion Process," *Journal of Economic Theory*, 49:33-83.
- CROW, E. L. AND K. SHIMIZU, editors (1988): *Lognormal Distributions: Theory and Applications*, New York and Basel: Marcel Dekker, INC.
- DUFFIE, D. (1996): *Dynamic Asset Pricing Theory*, Princeton, New Jersey: Princeton University Press.
- DUMAS, B., P. L. JENNERGREN, AND B. NÆSLUND (1995): "Siegel's Paradox and the Pricing of Currency options," *Journal of International Money and Finance*, 14:213-223.
- EAKER, M. AND D. GRANT (1985): "Optimal Hedging of Uncertain and Long-Term Foreign Exchange Exposure," *Journal of Banking and Finance*, 9:221-231.
- FEIGER, G. AND B. JACQUILLAT (1979): "Currency Option Bonds, Puts and Calls on Spot Exchange and the Hedging of Contingent Foreign Earnings," *Journal of Finance*, 34:1129-1139.
- FRANKE, G., R. C. STAPLETON, AND M. G. SUBRAHMANYAM (1998): "Who Buys and Who Sells Options: The Role of Options in an Economy with Background Risk," *Journal of Economic Theory*, 82:89-109.
- GARMAN, M. AND S. KOHLHAGEN (1983): "Foreign Currency Option Values," *Journal of International Money and Finance*, 2:213-237.

- GEMAN, H., N. EL KAROUI, AND J.-C. ROCHET (1995): "Changes of Numeraire, Changes of probability measure and Option Pricing," *Journal of Applied Probability*, 32:443–458.
- GIDDY, I. (1985): "The Foreign Exchange Option as a Hedging Tool," in Lessard, D., editor, *International Financial Management: Theory and Application*. New York: John Wiley.
- GRINOLD, R. C. AND R. N. KAHN (1995): *Active Portfolio Management*, New York: The McGraw-Hill Companies, Inc.
- GROSSMAN, S. J. AND J.-L. VILA (1989): "Portfolio Insurance in a Complete Market: A Note," *Journal of Business*, 62:473–476.
- HARRISON, J. AND D. KREPS (1979): "Martingales and Arbitrage in Multiperiod securities markets," *Journal of Economic Theory*, 20:381–408.
- HARRISON, J. AND S. PLISKA (1981): "Martingales and Stochastic Integrals in the Theory of Continuous Trading," *Stochastic Processes and Their Applications*, 15:313–316.
- HODRICK, R. (1987): "The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets," in Branson, W., editor, *Fundamentals of Pure and Applied Economics*. London: Harwood Academic Publishers.
- HUANG, C.-F. AND R. H. LITZENBERGER (1988): *Foundations for Financial Economics*, North Holland, Elsevier Science Publishing Co., Inc., New York.
- HULL, J. C. (2000): *Options, Futures, and Other Derivatives*, Prentice Hall International, Inc., New Jersey, fourth edition.
- KARATZAS, I. (1997): *Lectures on the Mathematics of Finance*, Rhode Island: American Mathematical Society.
- KERKVLiet, J. AND M. H. MOFFET (1991): "The Hedging of an Uncertain Future Foreign Currency Cash Flow," *Journal of Financial and Quantitative Analysis*, 26:565–578.
- LAPAN, H., G. MOSCHINI, AND S. D. HANSON (1991): "Production, Hedging, and Speculative Decisions with Options and Futures Markets," *American Journal of Agricultural Economics*, 20000:66–74.

- LEHRBASS, F. (1994): "Optimal Hedging with Currency Forwards, Calls, and Calls on Forwards for the Competitive Exporting Firm Facing Exchange Rate Uncertainty," *Journal of Economics*, 59:51-70.
- LELAND, H. (1980): "Who Should Buy Portfolio Insurance?," *Journal of Finance*, 35:581-594.
- LEWIS, K. (1995): "Puzzles in International Financial Markets," in Grossman, G. and K. Rogoff, editors, *Handbook of International Economics*, vol. III. New York: Elsevier Science B.V.
- MARKOWITZ, H. (1952): "Portfolio Selection," *Journal of Finance*, 7:77-91.
- MCCULLOCH, J. H. (1975): "Operational Aspects of the Siegel Paradox," *Quarterly Journal of Economics*, 89:170-172.
- MCDONALD, R. AND D. SIEGEL (1986): "The Value of Waiting to Invest," *Quarterly Journal of Economics*, 101:707-727.
- McKEAN, H. P. (1970): *Appendix: A Free Boundary Problem for the Heat Equation Arising from a Problem in Mathematical Economics*, Volume 3 of Merton (1970).
- MERTON, R. (1973): "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4:141-183.
- MERTON, R. C., editor (1970): *The collected Scientific Papers of Paul A. Samuelson*, MIT Press, Cambridge, Mass.
- (1971): "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3:373-413.
- MOSCHINI, G. AND H. LAPAN (1992): "Hedging Price Risk with Options and Futures for the Competitive Firm with Production Flexibility," *International Economic Review*, 33:607-618.
- (1995): "The Hedging Role of Options and Futures under Joint Price, Basis, and Production Risk," *International Economic Review*, 36:1025-1049.
- MUSIELA, M. AND M. RUTKOWSKI (1997): *Martingale Methods in Financial Modelling*, Berlin: Springer-Verlag.

NEWBERY, D. M. (1989): "Futures Markets, Hedging and Speculation," in Eatwell, J., M. Milgate, and P. Newman, editors, *The New Palgrave: Finance*. The Macmillan Press Limited.

ØKSENDAL, B. (1998): *Stochastic Differential Equations*, Berlin: Springer-Verlag, fifth edition.

OLSEN, T. AND G. STENSLAND (1992): "On Optimal Timing of Investment when Cost Components are Additive and Follow Geometric Diffusions," *Journal of Economic Dynamics and Control*, 16:39–51.

PELSSER, A. (2000): *Efficient Methods for Valuing Interest Rate Derivatives*, London: Springer-Verlag.

PERSSON, S.-A. AND T. TROVIK (1997): "Optimal Hedging Revisited: The Importance of a Risk Premium," Working Paper, Norwegian School of Economics And Business Administration, Department of Finance and Management Science.

REINER, E. (1992): "Quanto Mechanics," *Risk*, (5):59–63.

ROPER, D. E. (1975): "The Role of Expected Value Analysis for Speculative Decisions in the Forward Currency Market," *Quarterly Journal of Economics*, 89:157–169.

RUBINSTEIN, M. (1976): "The valuation of Uncertain Income Streams and the Pricing of Options," *Bell Journal of Economics*, 7:407–425.

SAMUELSON, P. (1965): "Rational Theory of Warrant Pricing," *Industrial Management Review*, 6:13–31.

SIEGEL, J. J. (1972): "Risk, Interest Rates, and the Forward Exchange," *Quarterly Journal of Economics*, 82:303–309.

————— (1975): "Reply - Risk, Interest Rates, and the Forward Exchange," *Quarterly Journal of Economics*, 88:173–175.

STAPLETON, R. C. AND M. G. SUBRAHMANYAM (1984): "The Valuation of Multivariate Contingent Claims in Discrete Time Models," *The Journal of Finance*, 39(1):207–228.

STEIL, B. (1993): "Currency Options and the Optimal Hedging of Contingent Foreign Exchange Exposure," *Economica*, 60:413–431.

STULZ, R. (1984): "Optimal Hedging Policies," *Journal of Financial and Quantitative Analysis*, 19:127-140.

ZHANG, P. (1997): *Exotic Options*, Singapore: World Scientific Publishing Co. Pte. Ltd.