

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Syrdal, Stig Arild

Working Paper A Study of Implied Risk-Neutral Density Functions in the Norwegian Option Market

Working Paper, No. 2002/13

Provided in Cooperation with: Norges Bank, Oslo

Suggested Citation: Syrdal, Stig Arild (2002) : A Study of Implied Risk-Neutral Density Functions in the Norwegian Option Market, Working Paper, No. 2002/13, ISBN 82-7553-204-3, Norges Bank, Oslo, https://hdl.handle.net/11250/2498655

This Version is available at: https://hdl.handle.net/10419/209812

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



ND https://creativecommons.org/licenses/by-nc-nd/4.0/deed.no





ANO 2002/13

Oslo December 9, 2002

Working Paper Securities Markets and International Finance Department

A Study of Implied Risk-Neutral Density Functions in the Norwegian Option Market

by

Stig Arild Syrdal

NORGES BANK

Working papers fra Norges Bank kan bestilles over e-post: posten@norges-bank.no eller ved henvendelse til: Norges Bank, Abonnementsservice Postboks 1179 Sentrum 0107 Oslo Telefon 22 31 63 83, Telefaks 22 41 31 05

Fra 1999 og senere er publikasjonene tilgjengelige som pdf-filer på www.norges-bank.no, under "Publisert".

Working papers inneholder forskningsarbeider og utredninger som vanligvis ikke har fått sin endelige form. Hensikten er blant annet at forfatteren kan motta kommentarer fra kolleger og andre interesserte.

Synspunkter og konklusjoner i arbeidene står for forfatternes regning.

Working papers from Norges Bank can be ordered by e-mail: posten@norges-bank.no or from Norges Bank, Subscription service, P.O.Box. 1179 Sentrum N-0107 Oslo, Norway. Tel. +47 22 31 63 83, Fax. +47 22 41 31 05

Working papers from 1999 onwards are available as pdf-files on the bank's web site: www.norges-bank.no, under "Published".

Norges Bank's working papers present research projects and reports (not usually in their final form) and are intended inter alia to enable the author to benefit from the comments of colleagues and other interested parties.

Views and conclusions expressed in working papers are the responsibility of the authors alone.

ISSN 0801-2504 ISBN 82-7553-204-3

A Study of Implied Risk-Neutral Density Functions in the Norwegian Option Market

Stig Arild Syrdal December 2002

Abstract

Option prices are assumed to contain unique information about how market participants assess the likelihood of different outcomes for future market prices. The main object of this study is to analyse the potential value of information contained in prices of options on the OBX index at Oslo Stock Exchange. The information is extracted using implied risk-neutral density functions. The study shows that there is a high level of uncertainty surrounding the implied density functions extracted from OBX options. Uncertainty introduced by using an average of the closing bid and ask quotation as a proxy for the option price, and the small range of actively traded strike prices, suggest that we should not place to much confidence in estimates sensitive to the tails of the implied density functions. The small range of actively traded strike prices is probably also a major reason for the differences often observed between various estimation techniques. Using information contained in OBX option prices in forecasting future market prices seems to be worthless. Some information about future volatility may be obtained, but not about the direction of future outcomes.

Keywords: Implied risk-neutral density functions, option pricing, market expectations. **JEL classification codes:** C10, G13.

Preface

This work has been carried out during my student internship in the Securities Markets and International Finance Department in Norges Bank in the period from January through May 2002, and is submitted in partial fulfillment of the requirements for the degree of Cand. Merc at the Norwegian School of Economics and Business Administration (NHH). Great thanks to Arild Lund, Sindre Weme, Bjørne-Dyre Syversten, Jon Bergundhaugen and especially Ketil Johan Rakkestad for the hospitality and backing I experienced during my stay. I would particularly like to thank Bernt Arne Ødegaard in the Research Department of Norges Bank, and my supervisor at NHH, Professor Steinar Ekern, for great support throughout this study. I am also grateful for useful discussions with Qaisar Farooq Akram, Magnus Andersson, Tom Bernhardsen, Knut Eeg, Øistein Røisland and Bent Vale. Thanks to Kari Hovde and Preben Danielsen at Oslo Stock Exchange for providing the necessary data.

Contents

1	Intr	Introduction			
2	Option Theory				
	2.1	Options	3		
	2.2	The Black-Scholes Model	4		
		2.2.1 Price Process and Distributional Properties	5		
		2.2.2 The Black-Scholes-Merton Differential Equation	6		
		2.2.3 Pricing Options with the Black-Scholes Model	7		
		2.2.4 Valuing Futures Options	9		
	2.3	The Volatility Smile	10		
3	Implied Risk Neutral Density Functions				
	3.1	General	13		
	3.2	The Risk Neutral Distribution			
	3.3	B Estimation of Risk Neutral Densities			
		3.3.1 Recovery of the Stochastic Process	15		
		3.3.2 The Risk-Neutral Valuation Equation	15		
		3.3.3 The Breeden-Litzenberger Result	18		
	3.4 Assessing the Uncertainty of Implied Risk-Neutral Density Functions .		20		
		3.4.1 Error Sources in Option Prices	21		
		3.4.2 Quantifying the Uncertainty of Implied Risk-Neutral Density Functions .	22		
	3.5	Information Content of Implied Risk-Neutral Density Functions	25		
4	Methodology				
	4.1	The Single Lognormal Model (SLN)	27		
	4.2	The Double Lognormal Model (DLN) $\ldots \ldots \ldots$	29		
	4.3	The Smoothed Implied Volatility Smile Method (SPLINE)	31		
		4.3.1 General Procedure	32		
		4.3.2 The Smoothing Spline Procedure	32		

		4.3.3	The Tails of the Distribution	36		
	4.4	Summ	ary Statistics for Implied Risk-Neutral Density Functions	38		
	4.5	Bid-A	sk Bounds	41		
	4.6	Testin	g the Information Content of Option Prices	42		
		4.6.1	Realized vs Implied Volatility	42		
		4.6.2	Information Content of Implied Skewness	43		
5	Ana	alyses		47		
	5.1	Data		47		
	5.2	Exam	ples of Implied Risk-Neutral Density Functions in the Norwegian Option			
		Marke	et	49		
		5.2.1	The Russian Crisis	50		
	5.3	Comp	aring Estimation Methods	54		
		5.3.1	Pricing Errors	54		
		5.3.2	Summary Statistics	57		
	5.4	Study	ing the Effect of the Bid-Ask Spreads	60		
		5.4.1	Stability of the Implied Distributions	60		
		5.4.2	Assessing Changes in the Market Sentiment	64		
	5.5	Are In	nplied Risk-Neutral Density Functions Useful as Leading Indicators? $% \mathcal{A} = \mathcal{A} = \mathcal{A}$	66		
		5.5.1	Studying Possible Leading Indicators	66		
		5.5.2	Realized versus Implied Volatility	69		
		5.5.3	Information Content of Asymmetries within the Implied Risk-Neutral Den-			
			sity Functions	71		
		5.5.4	Discussion of Results	73		
6	Sun	nmary	and Conclusions	75		
Bi	Bibliography					
\mathbf{A}	Appendices					

Chapter 1

Introduction

Market participants and policy-makers working in the financial markets use information embedded in prices on financial assets to analyse economic and financial development. In recent years, there has been a remarkable growth in the derivative markets and products such as futures and options are gaining increased popularity. The derivative markets provide a rich source of information for gauging market sentiment. Option prices are especially useful for extracting such information. Since an option's payoff depends on the future development of the underlying asset, the prices of different option contracts reflect the market participants' view of the likelihood that the contract will yield a positive payoff. Thus, by studying prices of options on a particular asset with different strike prices but with the same time to maturity, we may learn something about the probability that the market attaches to the asset being within a range of possible prices at some future date. A popular way of gaining this information is by estimating the so called *implied risk-neutral density function*. Under given assumptions, this function can be interpreted as the market's aggregate probability distribution for the price of the underlying asset at maturity. Thus, it may contain valuable information about the market participants' expectations regarding the future development of the underlying asset.

For example, the implied risk-neutral density function may tell us whether market participants place relatively greater probability on an upward price movement than on a downward movement. This is illustrated by distribution (2) in Figure 1.1. The large right tail of the distribution suggests that the market agents are positive about the future development in the underlying asset. The implied risk-neutral density function may also tell us whether the market believes that extreme upward or downward movements are likely to occur. For example, if we observe an implied density function as distribution (1) in Figure 1.1 at one date, and as distribution (3) on a later date, it suggests that the market participants have become more worried of extreme movements in the price of the underlying asset.

Due to the seemingly unique information embedded in the implied risk-neutral density functions about how market participants assess the likelihood of different outcomes for future market



Figure 1.1: Examples of implied risk-neutral density functions

prices, the implied density functions are gaining increased attention among academics, traders, investors and central banks. A large number of techniques for estimating risk-neutral density functions have been proposed in the literature. In this study, the most popular methods for extracting implied risk-neutral density functions are implemented. The study is performed on equity options on the OBX index at Oslo Stock Exchange. One objective is to compare the relative performance of the various methods and to study the uncertainty surrounding the estimation of implied density functions in the Norwegian option market. A second objective is to analyse the potential value of the information embedded in OBX option prices. The aim is to gain a better understanding of whether properties of implied risk-neutral density functions can be used as leading indicators in the Norwegian stock market.

The report is organized as follows. To provide the necessary theoretical foundation, Chapter 2 starts with a brief description of option contracts, and then discusses important characteristics of the most widely used model for valuing option contracts, the Black-Scholes model. Chapter 3 discusses various methods for extracting implied risk-neutral density functions, and review some of the earlier literature on this subject. A comprehensive description of the methodology applied in the present study is given in Chapter 4. The analyses are presented in Chapter 5. I first give examples of implied risk-neutral density functions during a financial stress event to illustrate how the shape of the distributions may change in response to such events. A comparison of the performance of the various estimation methods applied in the study is presented next. I then illustrate some of the uncertainty related to the estimation of implied distributions in the Norwegian option market. As an extension of this analysis, I also show how to take account of this uncertainty when assessing changes in the market sentiment. Finally, I asses the possibility of using properties of implied risk-neutral density functions as leading indicators in the Norwegian stock market. Summary and conclusions are given in Chapter 6.

Chapter 2

Option Theory

To provide the necessary theoretical foundation, I will in this chapter discuss basic elements of option theory. The chapter starts with a brief description of option contracts. I then present the most widely used model for valuing option contracts, the Black-Scholes model, and discuss important characteristics of the model.

2.1 Options

There are two basic types of options, *call* options and *put* options. A call option gives the holder the right to buy the underlying asset by a certain date for a certain price. A put option gives the holder the right to sell the underlying asset by a certain date for a certain price. Note that the holder is not obliged to exercise this right. The underlying assets include stocks, stock indices, foreign currencies, debt instruments, commodities, and futures contracts. The price at which the underlying asset can be sold or bought is called the *exercise price* or *strike price*. The date in the option contract is known as the *expiration date* or *maturity*. There are further two basic types of call and put options, so called American and European options. American options can be exercised at any time up to the expiration date, while European options can only be exercised at the expiration date itself. American options are most common. It is generally easier to analyse European options than American options, and some of the properties of the American options are therefore often deduced from properties of its European counterpart.

If the price of the underlying asset is above the exercise price, the holder of a call option can buy the underlying asset for a price lower than the market price. Hence, the holder of the option will obtain a positive cash flow if it is exercised immediately. The option is then referred to as being in the money. If this cash flow is bigger than the initial price paid for the option, the holder will earn a positive profit. If the price of the underlying asset is lower than the exercise price, the call option is referred to as being out of the money. The option will then give a negative cash flow if exercised immediately. Clearly, an out-of-the-money option will never be exercised immediately. For put options the situation is reversed. A put option is in the money if the price of the underlying asset is lower than the exercise price. The holder of the put option can then sell the underlying asset for a price above the market price, and receive a positive cash flow. If the price of the underlying asset is higher than the exercise price, the put option is out of the money, and will not be exercised. If the price of the underlying asset is equal to the exercise price an option is referred to as being at the money.

Every option contract has two sides. The trader who has taken the long position (i.e., has bought the option) is on one side, and the trader who has taken the short position (i.e., has sold or written the option) is on the on the other side. Hence, four basic option positions are possible; a long or short position in the call option, and a long or short position in the put option. By expressing the option position in terms of the payoff at maturity, excluding the initial price of the option, the four basic positions can be written as:

- 1. Long position in a call option: $\max(S_T-X, 0)$
- 2. Long position in a put option: $\max(X-S_T, 0)$
- 3. Short position in a call option: $-\max(S_T-X, 0) = \min(X-S_T, 0)$
- 4. Short position in a put option: $-\max(X-S_T, 0) = \min(S_T-X, 0)$

 S_T is the price of the underlying asset at maturity, and X is the exercise price. The four basic positions are illustrated in Figure 2.1.

2.2 The Black-Scholes Model

The breakthrough in option pricing came in the early seventies. Fisher Black, Myron Scholes and Robert Merton developed what we today know as the Black-Scholes model (Black and Scholes [1973], Merton [1973]). This model is today the most popular model for valuing European call and put options on non-dividend paying stocks. A short review of this model is presented next. The presentation is based on Hull [2000], Chapters 10-12.



Figure 2.1: Payoff from the four basic option positions

2.2.1 Price Process and Distributional Properties

The starting point of the Black-Scholes model is the assumption that the stock price follows a geometric Brownian motion. In discrete-time the stock price behavior can be expressed as:

$$\frac{\Delta S}{S} = \mu \Delta t + \sigma \varepsilon \sqrt{\Delta t} \tag{2.1}$$

The variable ΔS is the change in the stock price S in a small interval of time Δt . ε is a random variable drawn from a standardized normal distribution ($\varepsilon \sim N(0, 1)$). The parameter μ is the expected rate of return per unit of time from the stock, and σ is the volatility of the stock price. Both μ and σ are assumed constant. The model implies that the return of the stock can be expressed as the sum of a deterministic component, equal to the expected rate of return, and a stochastic component. Hence, the return is normally distributed with mean $\mu \Delta t$ and variance $\sigma^2 \Delta t$:

$$\frac{\Delta S}{S} \sim N(\mu \Delta t, \sigma^2 \Delta t) \tag{2.2}$$

By expressing (2.1) in continuous time $(\Delta t \to 0)$, the stock price dynamics becomes:

$$\frac{dS}{S} = \mu dt + \sigma dz \tag{2.3}$$

where the variable dz is a continuous Wiener process and is equal to $\varepsilon\sqrt{dt}$. Applying a mathematical result known as Ito's lemma¹ to equation (2.3), it can be shown that the price, f, of an option or another derivative written on the underlying stock S, has to satisfy the following relation²:

$$df = \left(\frac{\partial f}{\partial S}\mu S + \frac{\partial f}{\partial t} + \frac{1}{2}\frac{\partial^2 f}{\partial S^2}\sigma^2 S^2\right)dt + \frac{\partial f}{\partial S}\sigma Sdz$$
(2.4)

By comparing equation (2.3) and (2.4) we see that both S and f are affected by the same source of uncertainty, dz. This is a very important result in the derivation of the Black-Scholes model.

If we assume that the price of a non-dividend paying stock follows a geometric Brownian motion, it can be easily verified from equation (2.4) that the logarithm of the stock price must satisfy³:

$$dlnS = \left(\mu - \frac{\sigma^2}{2}\right)dt + \sigma dz \tag{2.5}$$

This implies that the logarithm of the stock price is normally distributed:

$$lnS_T - lnS_0 = ln\frac{S_T}{S_0} \sim N\left[\left(\mu - \frac{\sigma^2}{2}\right)T, \ \sigma^2 T\right]$$
(2.6)

$$\Rightarrow lnS_T \sim N\left[lnS_0 + (\mu - \frac{\sigma^2}{2})T, \ \sigma^2 T\right]$$
(2.7)

To sum up, if the stock price follows a geometric Brownian motion then the stock price is lognormally distributed and the return of the stock is normally distributed with constant variance.

2.2.2 The Black-Scholes-Merton Differential Equation

Equation (2.4) is the starting point for deriving the famous Black-Scholes-Merton differential equation. By combining the stock and the derivative in the same portfolio, the stochastic component $dz = \varepsilon \sqrt{dt}$ can be eliminated, making the portfolio riskless. This is always possible since the stock and its derivative are affected by the same sources of risk. The riskless portfolio can be obtained by going short in one derivative and long in an amount of $\frac{\partial f}{\partial S}$ shares. To eliminate the possibility for arbitrageurs to make riskless profit this portfolio must instantaneously earn the

¹See for example Hull [2000], Appendix 10A.

²See Hull [2000], page 229-231 for details.

³Can be verified by substituting f = lnS.

same rate of return as other short-term risk-free securities. Hence, the return of this portfolio must equal the risk-free interest rate, r. This gives us the Black-Scholes-Merton differential equation⁴:

$$\frac{\partial f}{\partial t} + rS\frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} = rf$$
(2.8)

It is important to realize that the portfolio used to derive equation (2.8) only is riskless instantaneously. When S and t change, also $\frac{\partial f}{\partial S}$ will change. For the portfolio to stay riskless it must be continuously rebalanced.

The Black-Scholes-Merton differential equation can be used to find the price of many different types of derivatives with the price S of a non-dividend paying stock as the underlying variable. The solution depends on the boundary conditions for the particular derivative. For example, for an European call option the key boundary condition is:

$$f = max(S - X, 0) \qquad when \ t = T \tag{2.9}$$

Equivalent for an European put option:

$$f = max(X - S, 0)$$
 when $t = T$ (2.10)

T is the time at maturity.

2.2.3 Pricing Options with the Black-Scholes Model

If we solve the Black-Scholes-Merton differential equation (2.8) with the proper boundary conditions, we get the Black-Scholes model for pricing European call and put options on non-dividend paying stocks:

$$c = S_0 N(d_1) - X e^{-r\tau} N(d_2)$$
(2.11)

$$p = Xe^{-r\tau}N(-d_2) - S_0N(-d_1)$$
(2.12)

where

$$d_{1} = \frac{\ln(\frac{S_{0}}{X}) + (r + \frac{\sigma^{2}}{2})\tau}{\sigma\sqrt{\tau}}$$
(2.13)

$$d_2 = d_1 - \sigma \sqrt{\tau} \tag{2.14}$$

⁴See Hull [2000] Chapter 11, page 246-248, for more details.

c is the price of an European call option, p is the price of an European put option, X is the exercise price, S_0 is the stock price today, τ is the remaining time to maturity⁵, and N() is the standard cumulative normal distribution function.

The Black-Scholes pricing formula can also be derived using a principle known as riskneutral valuation. This is a result of a very important property of the Black-Scholes-Merton differential equation. None of the parameters in the equation are affected by the risk preferences of investors. In other words, pricing of derivatives with the Black-Scholes-Merton differential equation is independent of risk preferences, and we can assume that all investors act as if they are risk-neutral. Thus, when valuing an option we calculate the expected payoff assuming that the expected return from the underlying asset is the risk-free interest rate, and use the same risk-free interest rate to discount the expected payoff. For example, the price of an European call and put option can be written as:

$$c = e^{-r\tau} \hat{E}[max(S_T - X, 0)]$$
 (2.15)

$$p = e^{-r\tau} \hat{E}[max(X - S_T, 0)]$$
(2.16)

where $\hat{E}()$ represents the expectation taken with respect to a risk-neutral distribution with expected return equal to the risk-free interest rate. Assuming that the stock price follows a geometric Brownian motion, the corresponding risk-neutral distribution is lognormal. The Black-Scholes model can then be derived using equation (2.15) and (2.16)⁶.

The Black-Scholes model can be easily modified to take account of dividends. Dividends have the effect of reducing the stock price on the ex-dividend date. Hence, by assuming that the amount and timing of the dividends during the life of an option can be predicted with certainty, the stock price on the ex-dividend date can be adjusted to take account of the dividends. To simplify the analysis it is generally assumed that the stock pays a continuous dividend yield at a rate δ per year. The continuous dividend yield causes the growth rate in the stock price to be reduced by an amount δ . Thus, when valuing an European option with remaining time to maturity τ , we reduce the current stock price from S_0 to $S_0 e^{-\delta \tau}$ and the expected return from r to $r - \delta$, and value the option as though it pays no dividends.

In addition to the assumptions that the stock price follows a geometric Brownian motion with constant mean and volatility, no dividends during the life of the option and no riskless arbitrage opportunities, the derivation of the general Black-Scholes model also assumes no restrictions on short sale, no taxes or transaction costs, continuous security trading, that all securities are perfectly divisible, and that the risk-free interest rate is constant and the same for all maturities.

⁵Throughout this study, T is the time at maturity, while τ is the remaining time to maturity.

⁶See Hull [2000] appendix 11A for details.

For a discussion on relaxing the assumptions in the Black-Scholes model, see Hull [2000] Section 17.6.

2.2.4 Valuing Futures Options

The data used in the present study consist of options on the OBX index at Oslo Stock Exchange. Also futures contracts are traded on the OBX index. These are agreements to buy or sell the OBX index at a certain future time for a certain price. The price of these contracts is called the futures price (F). To eliminate riskless arbitrage opportunities, the futures price must equal⁷:

$$F_0 = S_0 e^{r\tau} \tag{2.17}$$

If $F_0 > S_0 e^{r\tau}$, riskless profit can be obtained by shorting the futures contract and buying the asset. Similarly, if $F_0 < S_0 e^{r\tau}$, riskless profit can be obtained by shorting the asset and buying the futures contract.

The OBX future matures at the same time as the OBX option. At maturity, the price of a futures contract must be equal to the spot price of the underlying asset, $F_T = S_T$. If we consider European options, this implies that the futures option and the spot option are depending on the same underlying variable. Consequently, options on futures and options on the spot with the same strike price and time to maturity are in theory equivalent. This means that futures contracts on the OBX index can be used as a proxy for the underlying assets in the OBX index. OBX options are therefore priced as though they are options on the index futures.

If we assume that futures prices have the same lognormal property as assumed earlier, European futures options can be valued by extending the general Black-Scholes model. Black [1976] shows that the call price, c, and the put price, p, of European futures options can be valued by substituting S_0 with $F_0 e^{-r\tau}$ in (2.11) and (2.12):

$$c = e^{-r\tau} [F_0 N(d_1) - X N(d_2)]$$
(2.18)

$$p = e^{-r\tau} [XN(-d_2) - F_0N(-d_1)]$$
(2.19)

where d_1 and d_2 now are given by:

$$d_1 = \frac{\ln(\frac{F_0}{X}) + \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}}$$
(2.20)

$$d_2 = d_1 - \sigma \sqrt{\tau} \tag{2.21}$$

⁷Assuming that no dividends are paid.

Note that σ is the volatility of the futures price. Black's model (often termed the Black-76 model) does not require that that the option contract and the futures contract matures at the same time.

2.3 The Volatility Smile

In the Black-Scholes framework the option price is a function of five (or six) variables; the current stock price S, the exercise price X, the risk-free interest rate r, time to maturity τ and the volatility σ (and the dividend rate δ). If we are neglecting dividends, all the variables, except the volatility, are variables that can be directly observed when the option is priced. The price of an option is therefore depending on the market's opinion about the future volatility of the underlying asset upon which the option is written. Consequently, the volatility parameter is the single most important parameter when valuing options.

The volatility that makes the theoretical option price calculated from the Black-Scholes model equal to the observed option price, is called the implied (Black-Scholes) volatility. The implied volatility can be easily found by an iterative search procedure using the Black-Scholes formula. Volatility, or standard deviation, is often used as a measure of risk. By calculating the implied volatility of an option we obtain a point estimate of the risk that the market assigns to the underlying asset in the next period. Hence, the implied volatility contains useful information about the market participants' belief about the future volatility of the underlying asset.

According to the Black-Scholes model, implied volatilities from options should be the same regardless of which option is used to compute the volatility. In practice, this is usually not the case⁸. Options on the same underlying asset with different strike prices and maturities yield different implied volatilities. The pattern of the Black-Scholes implied volatilities with respect to strike prices has become known as the volatility smile. A typical shape of a volatility smile for an equity options is illustrated in Figure 2.2.

The volatility smile for equity options is sometimes referred to as a "volatility skew" because typically the implied volatility decreases as the strike price increases. This means that outof-the money puts and in-the-money calls have a greater implied volatility than in-the-money puts and out-of-the-money calls of equivalent maturity. The existence of a volatility smile is clearly inconsistent with the Black-Scholes model. If only one volatility is used to price options with different strikes, the pricing errors will be systematically related to the strike price. It has also been shown that the smile depends on the options' maturities. The inconsistency of the Black-Scholes model means that options are not priced as though the underlying asset follows a geometric Brownian motion, and that the underlying asset price is lognormally distributed.

Rubinstein [1994] points out that the smile effect has become consistently pronounced after

 $^{^8 \}mathrm{See}$ for example Rubinstein [1994].



Figure 2.2: A typical volatility smile for an equity option

the stock market crash of 1987. Earlier tests of option pricing were more or less supportive of the Black-Scholes model. Rubinstein calls this phenomenon "crash-o-fobia". His explanation is that after 1987, traders have become intensely concerned about the possibility of similar crashes. Thus, the cost of crash protection, represented by out-of-the-money puts, has increased. Another explanation is the leverage effect. When the stock price declines, the company's equity decline and the leverage increases. The equity becomes more risky and the volatility increase.

The shape of the volatility smile can be used to extract information about the market's aggregate opinion about the underlying distribution, and how it differs from the simple lognormal distribution⁹. The shape of the distribution implied by the market is directly related to the slope and the convexity of the smile curve. A negative (positive) slope indicates that the implied market distribution is skewed to the left (right). The amount of skewness depends on the size of the slope. A negative volatility skew and the corresponding implied distribution are illustrated in Figure 2.3.



Figure 2.3: A skewed volatility smile and the corresponding implied distribution

⁹See Hull [2000], page 438-440.

The convexity of the volatility curve is related to the fatness of the tails. A convex curve implies that the corresponding underlying distribution has fatter tails than the lognormal distribution, i.e. the implied distribution is more leptokurtic. This is illustrated in Figure 2.4.



Figure 2.4: A convex volatility smile and the corresponding implied distribution

As pointed out in this section, options are generally not priced as though the distribution of the underlying asset is lognormally distributed. Practitioners use different estimates of stock return volatility to value different options. In the next chapter we shall see how to use information contained in option prices to derive the distribution of the underlying asset implied by the market participants.

Chapter 3

Implied Risk Neutral Density Functions

3.1 General

The volatility implied from option prices contains useful information about the market participants' belief about the risk associated with the underlying asset in the future. But more information can be extracted from option prices. Since an option's payoff depends on the future development in the underlying asset, the prices of different option contracts reflect the market participants' view of the likelihood that the contract will yield a positive payoff. Thus, by studying prices of options on a particular asset with different strike prices but with the same time to maturity, they may tell us something about the probability that the market attaches to the asset being within a range of possible prices at some future date.

A popular way of gaining this information from option prices is by estimating the so called implied risk-neutral density function. Under given assumptions, this function can be interpreted as the market's aggregate probability distribution for the price of the underlying asset at maturity. It may therefore contain valuable information about market expectations. For example, the implied risk-neutral density function may tell us whether market participants place relatively greater probability on a downward price movement than on an upward movement. Or, whether they believe that extreme upward and downward movements are likely to occur. This is information that can not be extracted from the lognormal property of the Black-Scholes model. Generally, the implied distribution differs from the lognormal density function underlying the Black-Scholes model.

3.2 The Risk Neutral Distribution

The density function implied from option prices is strongly related to the definition of statecontingent prices. State-contingent prices are the prices of securities that promise one unit of money if a certain state occurs at a given date, and zero otherwise. If you hold a statecontingent claim for all possible states for a given date, you will receive one unit of money for certain. Hence, the sum of the state-contingent prices across all states has to equal the price of a risk-free zero coupon bond that pays one unit of money at maturity for sure, i.e. one discounted by the risk-free interest rate. By normalizing the state contingent prices by the inverse of the price of the bond, they will sum up to one. These normalized state-contingent prices represent the risk-neutral densities for the different states. The densities are risk-neutral simply because the payoff from a state-contingent claim for a certain state is riskless. In a complete market¹ it will be possible to recover the complete risk-neutral density function. Ross [1976] proves this by using a set of European option prices. An exact expression for the relationship between European option prices and the risk-neutral probability distribution was derived a couple of years later by Breeden and Litzenberger [1978]:

$$\frac{\partial^2 c(X,\tau)}{\partial X^2} = e^{-r\tau} q(S_T) \tag{3.1}$$

Equation (3.1) tells us that the second derivative of an European call price function (c) taken with respect to its strike price (X) is equal to the continuously discounted risk-neutral distribution (q) for the price of the underlying asset at maturity (S_T) . τ is the remaining time to maturity. A further discussion of equation (3.1) is given in Section 3.3.3.

From the above discussion it should be clear that the probability distribution implied from option prices actually is a distribution of normalized state-contingent prices, i.e. risk-neutral densities. This should not be surprising. As discussed in the previous chapter are prices of options and other derivatives independent of the investor's degree of risk aversion. Consequently, these prices can not give any information about risk preferences.

The risk-neutral densities are equivalent to the "true" market densities only if there is no aggregate risk in the market or assuming risk-neutrality. If we assume that there is aggregate risk in the market and allow for risk aversion, the two distributions will naturally differ. A state that may have a relatively high probability in the risk neutral-density may have a relatively low statistical probability of actual occurring, but the market value a unit of wealth much higher in this state. It is thus difficult to distinguish between changes in the "true" probabilities and changes in the value of wealth in the different states. But even if the probability distribution derived from option prices is risk-neutral, it may still contain valuable information. Rubinstein

¹In a complete market it is either as many securities traded as there are states of the world, or as many dynamically rebalanced portfolios feasible as there are states of the world.

[1994] shows that if we assume that the representative investor has constant relative risk aversion, the "true" distribution will shift to the right, but the shape of the distribution is unchanged. From these results, it is reasonable to assume that changes in the implied density functions can give us valuable information about alterations in the market's opinion about the future development.

3.3 Estimation of Risk Neutral Densities

A large number of techniques for estimating risk neutral densities have been proposed in the literature in recent years. In this section I will only present the most general methods for extracting risk neutral probabilities. I recommend Jackwerth [1999] for a more comprehensive discussion of different estimation techniques. Notice that all the techniques presented here ignore any complications induced by the early exercise feature of American options, and can therefore only be applied to European options.

3.3.1 Recovery of the Stochastic Process

A general method for estimating implied risk-neutral densities is to first assume a particular stochastic price process for the underlying asset, and use observed option prices to recover the unknown parameters of the specified price process. The probability distribution is then derived from the stochastic price process. For example, in the Black-Scholes model, the assumption that the price process follows a geometric Brownian motion with a constant expected drift rate and constant volatility gives a lognormal risk-neutral density function for the underlying price. Malz [1996] uses this method to show that if the exchange rate evolves according to a special jump-diffusion process (jumps superimposed upon a geometric Brownian motion), the implied risk-neutral distribution is a mixture of two lognormal distributions. For more complex stochastic processes, the risk-neutral density can not be computed in closed-form and must be approximated by numerical methods.

3.3.2 The Risk-Neutral Valuation Equation

A simpler approach to derive the risk-neutral density function is to assume a particular parametric form for the implied distribution, and use observed option prices to recover the parameters of the distribution. A great advantage of this method compared to the previous one is that while a specified stochastic price process implies a unique distribution, is a given risk-neutral density function consistent with a range of different stochastic processes. As pointed out in the previous chapter, an option price can be derived using risk-neutral valuation. This implies that the price of an option can be expressed as the expected value of the option discounted by the risk-free interest rate, where the expected value is calculated from the risk-neutral density function. Hence, the price c of an European call option and the price p of an European put option at time t can be written as:

$$c = e^{-r\tau} \int_{0}^{\infty} q(S_T, \theta) max(S_T - X, 0) dS_T = e^{-r\tau} \int_{X}^{\infty} q(S_T, \theta)(S_T - X) dS_T$$
(3.2)

$$p = e^{-r\tau} \int_{0}^{\infty} q(S_T, \theta) max(X - S_T, 0) dS_T = e^{-r\tau} \int_{0}^{X} q(S_T, \theta)(X - S_T) dS_T$$
(3.3)

where r is the risk-free interest rate, τ is the remaining time to maturity, and $q(S_T, \theta)$ is the risk-neutral density function for the underlying asset at maturity T with parameter vector θ .

If we assume a particular form for the risk-neutral density function $q(S_T, \theta)$, the parameters of the distribution can be recovered by minimizing the squared deviation between observed option prices and theoretical option prices calculated from equation (3.2) and (3.3) across all exercise prices for a given maturity². Since call and put options are priced off the same underlying distribution, both sets of prices are included in the minimization problem. The minimization problem can then be written as:

$$\min_{\theta} \left[\sum_{i=1}^{m} (c_i - c_i^*)^2 + \sum_{j=1}^{n} (p_j - p_j^*)^2 \right]$$
(3.4)

where

 $c_i^*, p_j^* = \text{Observed option prices}$ $c_i, p_j = \text{Theoretical option prices calculated from (3.2) and (3.3) with density function <math>q(S_T, \theta)$

 θ = Vector of parameters for density function

In the absence of arbitrage opportunities, the forward price of the underlying asset must equal the mean of the implied risk-neutral density function. Generally, this relationship is included in the minimization problem in (3.4) by adding the squared deviation between the futures price and the mean of the distribution. Alternatively, this relationship can be imposed as a constraint in the minimization problem. A disadvantage of the latter approach is that the constraint usually will be binding and reduce the goodness-of-fit.

A key question is naturally which parametric form to assume for the risk-neutral density function. According to the Black-Scholes model, the distribution is lognormal. Then only two parameters need to be estimated, the mean and the volatility. Thus, the implied distribution can be obtained relatively easy. However, quite a few empirical studies have pointed out that

 $^{^2 \}mathrm{See}$ for example Bahra [1997].

prices on financial assets seldom are lognormally distributed³. Hence, a more flexible density function is required. A widely used method is to assume that the risk-neutral density function is the weighted sum of several independent lognormal distributions⁴. For example, Melick and Thomas [1997] extract the implied risk-neutral density function by assuming that the distribution is a weighted sum of three lognormal distributions. The study is performed on Americanstyle options on crude oil futures⁵ where the available range of strike prices is relatively large. Generally, options are traded across a smaller range of exercise prices, and therefore, the number of distributional parameters that can be estimated from the data is limited. Hence, it is more usual to use a mixture of two lognormal distributions (see Bahra [1997]). This form is sufficiently flexible to capture features such as skewness and fat tails that we might expect to find implicit in the data.

In the litterature, several other approximating functions for the implied distribution have been used. Examples of generalized distributions are the gamma and exponential distributions⁶ and different types of the Burr distribution⁷. Madan and Milne [1994] use a quite different approach to obtain the approximating function. They specify the normal distribution as a "prior" distribution and add correction terms to it by using a Hermite polynomial expansion. This method has been used in several other studies with different types of polynomial expansion. See Jackwerth [1999] for a more extensive description of these methods.

A somewhat similar approach is applied by Rubinstein [1994] and Jackwerth and Rubinstein [1996]. They take account of bid-ask bounds on the underlying asset price and the option prices, and use an optimization algorithm to find the implied risk-neutral density function, among all possible distributions which satisfy the required bid-ask constraints, that is "closest" to the prior lognormal distribution. Buchen and Kelly [1996] also use the lognormal density as a prior distribution, but apply a Bayesian maximum entropy approach to find the posterior distribution. This approach is in some respect is similar to the optimization method applied by Rubinstein [1994] and Jackwerth and Rubinstein [1996]. The main difference is related to the choice of objective function. While Rubinstein [1994] and Jackwerth and Rubinstein [1994] and Jackwerth and Kelly [1996] apply the cross-entropy function⁸. The rationale for the maximum entropy method is that a distribution that maximizes the entropy is least prejudiced with respect to unknown or missing information.

³See for example Campbell et al. [1997] page 16.

⁴The sum itself is not lognormally distributed.

⁵The early exercise feature of American options is solved by deriving bounds on the option price in terms of the terminal density function.

⁶See for example Aparicio and Hodges [1998].

⁷See for example Sherrick et al. [1996b], Sherrick et al. [1996a], Sherrick et al. [1992].

 $^{^{8}-\}sum p_{1}\ln[\frac{p_{1}}{p_{0}}], p_{0}$ is the prior distribution and p_{1} the posterior distribution.

3.3.3 The Breeden-Litzenberger Result

An alternative class of estimation techniques for recovering implied risk-neutral density function applies the Breeden-Litzenberger result in equation (3.1). Equation (3.1) tells us that the second derivative of an European call price function taken with respect to its strike price is equal to the continuously discounted risk-neutral density function, i.e. the normalized state price. In other words, if we can express the call prices as a function of the strike price, the implied distribution can be easily obtained by differentiation, either analytically or numerically. In the absence of arbitrage opportunities, $c(X, \tau)$ is convex and monotonically decreasing in X. This is a requirement for obtaining a positive density function. Equation (3.1) can be derived by looking at the relationship between state-contingent prices and option prices, or by differentiating the risk-neutral valuation equation directly.

As a first approximation to the implied risk-neutral density function we can generate so called risk-neutral histograms. If we start with the discrete version of equation (3.1), using a simple finite difference approximation for the second derivative, and solve for the density function we get:

$$q(S_T) = e^{r\tau} \frac{[c(S_T + \Delta S_T, \tau) - c(S_T, \tau)] - [c(S_T, \tau) - c(S_T - \Delta S_T, \tau)]}{(\Delta S_T)^2} |_{X = S_T}$$
(3.5)

Equation (3.5) is simply the normalized state-contingent price for $S_T = X$. From this expression we can estimate the approximate risk-neutral densities for the range of available strike prices, and the risk-neutral histogram can be established. A drawback of this method is that it requires that the options are traded at equally spaced strikes. In addition, due to the limited range of exercise prices traded in the marked, there is no systematic way of modeling the tails of the histogram. A further problem with this simple approximation is that it can not adjust for noise in the observed option prices. For example, if observed prices exhibit small but sudden changes in convexity or small degrees of concavity across strike prices, we may get spurious results.

To establish a continuously risk-neutral distribution function we need to apply interpolation techniques. One possibility is to interpolate the call option prices directly. This requires that the interpolated price function satisfies the monotonic and convexity constraints, and that the expression is twice differentiable. In the literature, both parametric and non-parametric methods are employed. The parametric approach imposes a particular parametric functional form directly on the observed option prices and estimates the functional parameters by minimizing the errors. For example, Bates [1991] fit a cubic spline ⁹ to the observed data. But there are several technical disadvantages of interpolating the call option pricing function directly. Small fitted

⁹A cubic spline consists of piecewise third order polynomials.

price errors may have large effects on the estimated risk neutral density, especially in the tails. Also, the form of the call price function may cause problems. Generally, the call price function has large curvature for options near at-the-money and little curvature for options far away from at-the-money. Consequently, a relatively large number of degrees of freedom are required to fit the function accurately.

Aït-Sahalia and Lo [1998] use a non-parametric method based on statistical kernel regressions to generate the relationship between the option price and the strike price. The non-parametric kernel estimator¹⁰ attempts to estimate the risk-neutral probabilities as a *fixed* function of certain economic variables. Thus, instead of just using a cross-section of data at a single point in time, they use a cross-sectional time-series to obtain the call price function. They estimate the call price function for two different sets of explaining variables. One set includes the stock price, strike price, time to expiration, interest rate, and dividend yield. The other set includes only the forward price, strike price, and time to expiration. Methods based on kernel regressions are extremely data-intensive and due to the limited range of available strike prices at a given point in time, they are generally hard to implement.

An alternative method for deriving implied risk-neutral densities based on the Breeden-Litzenberger result was proposed by Shimko [1993]. Instead of interpolating the call price function directly, he first interpolates the volatility smile, i.e. he interpolates the implied volatilities across strike prices. The volatility can then be written as a function of the strike price. By substituting this expression in the Black-Scholes model, the call price can be expressed as a continuous function of the strike price. By differentiating the call price function twice, the riskneutral density function can be extracted. Since the range of available strike prices is limited, the implied distribution will only expand between the lowest and highest strike price. Shimko solves this problem by fitting a lognormal distribution in each tail such that the total distribution sum up to one. Note that this method does not require that the Black-Scholes model is correct. The Black-Scholes model is simply used to transform the data from one space to another.

Today, variants of Shimko's method are widespread. Malz [1997] modifies Shimko's technique by interpolating the implied volatilities across deltas¹¹ instead of across strike prices. The advantage of this method is that in the implied volatility/delta space, options close to at the money are less grouped together than options far away from at the money. A given change in the strike price near the spot price gives a relatively large change in delta, while farther away from the money a corresponding change in strike price translates into a smaller change in delta. Hence, a greater "shape" near the centre of the distributions is permitted. Malz follow Shimko

¹⁰In one dimension, a kernel density estimator can be thought of as a way of smoothing a histogram. The smoothing is usually accomplished by constructing an assumed probability function around each data point. The overall density function is the weighted sum of the individual density functions.

¹¹The delta of an option is defined as the rate of change of the option price with respect to the price of the underlying asset.

in using a low-order polynomial as the smoothing function. Campa et al. [1998] on the other hand, interpolate in the implied volatility/strike price domain but introduce a new methodology for fitting the implied volatility curves. Instead of using a single polynomial they apply a cubic smoothing spline, i.e. a number of cubic polynomials joined together to a smooth curve. This permits the user to control the smoothness of the fitted function. Bliss and Panigirtzoglou [1999] combine the two methods. They follow Malz [1997] in interpolating in implied volatility/delta space and Campa et al. [1998] in using smoothing splines to fit the function.

Since the range of available (liquid) strike prices generally is limited to an area around at the money, the tails of the distribution represent a problem. Unlike Shimko [1993], Malz [1997] does not make any special assumptions for the tails. He allows the fitted curve to cover the entire range of possible deltas. Hence, the complete density function can be extracted. Bliss and Panigirtzoglou [1999] also let the curve span the entire implied volatility/delta space. They assume that the spline function is linear outside the region of observations. Campa et al. [1998], working in the implied volatility/strike price domain, use polynomials from the first and last region to extend the volatility smile left and right, respectively, and treat the smile as flat beyond that.

Several studies have compared how the different estimation methods perform relative to each other. As pointed out in Jackwerth [1999], the various methods generate rather similar risk-neutral densities unless we have very few option prices. Examples of such studies are Campa et al. [1998], Coutant et al. [2000], MacManus [1999] and Sherrick et al. [1996b]. In these studies the implied distributions are simply compared with respect to the various moments of the distributions and the in-sample goodness-of-fit. A evident weakness of such studies is that they do not consider the stability and robustness of the different estimation techniques. This is the subject of the next section.

3.4 Assessing the Uncertainty of Implied Risk-Neutral Density Functions

As emphasized in the previous section, a large number of techniques for estimating risk-neutral densities have been developed in the recent years. But relatively few studies on risk-neutral densities have focused on the uncertainty surrounding the estimated distributions. For example, how much confidence can we place in the summary statistics of the implied distribution, and how can we decide when changes in the implied distribution are due to alteration in market expectation, and not just noise? To answer these questions we need to quantify the measurement errors associated with the risk-neutral densities.

3.4.1 Error Sources in Option Prices

A first step in assessing the uncertainty of the estimated risk-neutral densities is to identify possible sources of measurements errors. An important source of error is related to the option prices used as input in the estimation. Possible errors are^{12} :

- Liquidity may be reflected in the option prices.
- Large bid-ask spreads.
- Only a narrow spectrum of strike prices is available.
- Non-synchronious quotes for the option price and the underlying asset.
- The sample may include strikes that have not been traded during the trading day.
- Data errors due to erroneous recording.
- Errors arising from quoting, trading and reporting prices in discrete increments.

A general problem in most derivative markets is the low liquidity for options being deep out of the money and deep in the money. The low liquidity of these options makes the prices less reliable and reduce the "accuracy" of the estimation. However, the problem can be avoided by only using the most liquid strikes when estimating the risk-neutral densities. But restricting the estimation in this way limits the range of available strikes, and makes the spectrum of strike prices even more narrow. Preferably, the range of strike prices should be as wide as possible. Option prices can only provide information about the underlying density at their respective strike prices. Consequently, outside the regions of available strikes, the distributions depend more on the choice of estimation technique than on the data. Thus, a more narrow spectrum of strike prices increases the uncertainty related to the tails of the distribution.

If real-time quotations are used, large bid-ask spreads may occur. This creates a problem regarding which price to use as input in the estimation, especially for out-of-the-money options. For these options, the bid-ask spreads become a higher percentage of the option premium, and may consequently lead to a misrepresentation of the underlying economic price.

If settlement prices are used instead of real time prices, the problem of large bid-ask spreads is avoided, but two major concerns arise. First, the market information used by the exchange when setting settlement prices at the end of the day is likely to be non-synchronous due to the infrequently trade of most option prices and the great variations in time-of-last-trade. However, this potential problem of non-synchronicity may be reduced by using only the most liquid strikes. Second, the sample may include strikes that have not been traded during the trading day. In

¹²Discussed in Bliss and Panigirtzoglou [1999], Melick and Thomas [1997], and Andersson and Lomakka [2001].

order to obtain more reliable density functions, these observations should be omitted from the analysis.

The problem of erroneous recording is more difficult to control. One way of reducing possible data errors is to screen the data for arbitrage opportunities. The fact that option prices are quoted for discrete strike prices is also a source of error. Even if the there are no other errors in the observed option prices, we cannot know to an accuracy of less than one half a tick at what price the option would have traded if prices were quoted on a continuous basis.

Since not all of these errors can be eliminated, there will always be some uncertainty attached to the estimated risk-neutral densities. Thus, to assess this uncertainty we need methods to quantify the measurement errors.

3.4.2 Quantifying the Uncertainty of Implied Risk-Neutral Density Functions

In the present study, I am going to study some of the uncertainty surrounding the implied risk-neutral density functions extracted from OBX options. I will therefore review some of the previous literature on this subject.

Söderlind and Svensson [1997] are the first in the literature to explicit consider the uncertainty related to the risk-neutral density estimation, and to derive a confidence band for the distribution. They assume that the correct model for the risk-neutral density is a mixture of two lognormal distributions, and that the actual prices differ from theoretical prices with a random error term. The parameters are obtained by minimizing the pricing errors in a non-linear least squares estimation. Consequently, the parameters of the distribution are approximately (and asymptotically) normally distributed. To account for heteroscedastic price errors they apply a heteroscedastic-consistent estimator of the covariance matrix. They then apply the delta method to obtain an approximate 95 percent point-by-point confidence band for the density function¹³.

Melick and Thomas [1998] also construct confidence bands by assuming that the risk-neutral density is a mixture of two lognormal distributions, but apply Monte Carlo simulations to derive them. They assume that the error between the estimated and the true parameter is multi-variate normally distributed with zero mean, and use the Hessian at the maximum likelihood solution as the estimated parameter variance-covariance matrix. From the assumed parameter distributions, new sets of parameters are randomly drawn, and new densities created. From these densities the confidence band is constructed.

For both Söderlind and Svensson [1997] and Melick and Thomas [1998], the 95 percent confidence bands appear to be quite narrow. This indicates that the uncertainty of the estimated risk-neutral densities is small. However, Melick and Thomas [1998] point out that the error terms are not independent, and thus, invalidate the results of the Monte Carlo method¹⁴. They

¹³Represents the confidence intervals for each single density point.

¹⁴Monte Carlo simulations relies on independent error terms in addition to the normality assumption.

therefore propose to use a bootstrap method to derive the confidence band. The idea behind a bootstrap method is to create a pseudo-sample by drawing with replacement from available observations, and then to estimate the model based on this sample. In this way, no structure is imposed on the error terms. By repeating this a large number of times, a set of parameters estimates is obtained and the pseudo-densities can be created. With the bootstrap method the confidence bands appear to be extremely wide. Melick and Thomas suggest that the bootstrap method is not capable of adequately quantifying the uncertainty, and that this is a result of the interdependence of probability measures derived from option prices with adjacent strike prices.

Instead of introducing disturbances in the *parameters* of the implied distribution, Söderlind [2000] proposes to perturb the fitted option *prices*. He starts out by estimating the implied risk-neutral density with the double lognormal method. He then adds error terms to the estimated theoretical option prices, and re-estimate the model using the simulated price set. This is repeated 100 times. The error terms are generated in two different ways; by drawing randomly from an i.i.d. normal distribution with the same variance as the original price errors, and by bootstrapping the original errors. Both methods produce relatively narrow confidence bands.

A somewhat similar approach is followed by Bliss and Panigirtzoglou [1999]. The aim of their study is to test the relative effects of measurements error on the stability of estimated riskneutral densities using the smoothed implied volatility smile method and the double lognormal method. To obtain simulated prices, they perturb the observed option prices by a random number uniformly distributed from minus and plus one half of the contracts "tick size". The risk-neutral densities are then calculated by the two methods for 100 simulated price-sets. Based on the accuracy and stability of the estimated summary statistics, Bliss and Panigirtzoglou [1999] conclude that the smoothed implied volatility smile method outperforms the doublelognormal method. Bliss and Panigirtzoglou do not explicitly calculate confidence bands for the distributions. Instead they calculate confidence intervals for the summary statistics. As they point out, the confidence intervals for higher order statistics, such as skewness, are sometimes so large that the estimates are useless.

A similar study is performed by Cooper [1999]. He evaluates the ability of the smoothed implied volatility smile method and the double lognormal method to recover simulated distributions based on the Heston [1993] stochastic volatility model. Cooper applies the same methodology as Bliss and Panigirtzoglou [1999]. The results of this study are also in favor of the smoothed implied volatility smile method.

Andersson and Lomakka [2001] suggest an extended (and improved) method for evaluating the robustness of implied risk-neutral densities. Since the error terms are unlikely to be normally distributed, using Monte Carlo simulation is not a valid method. Thus, they propose to use bootstrap methods in line with Söderlind [2000]. However, this method does not correct for possible heteroscedastic error terms. As Melick and Thomas point out, the error terms are not independent of either option type (call/put) or strike price. This error structure is also present in the OMX data used by Andersson and Lomakka. Thus, they aim at taking the heteroskedastic nature of the pricing errors into account by grouping the data in an appropriate manner.

Andersson and Lomakka apply two methods for deriving confidence bands; bootstrap from historical errors and bootstrap from actual error terms. To apply the first method they generate the historical patterns of the error terms. The double-lognormal method and the smoothed implied volatility method are estimated from January 1993 until June 2001 with 30 days to expiration. The error terms for the double lognormal method are then grouped according to the relative strike price¹⁵, a grouping strategy which consequently takes into account that the strike price range differs over time. For the smoothed implied volatility smile method the error terms are grouped over the range of delta¹⁶. The simulated prices are then created by drawing an error term with replacement from its corresponding group for each observation, and adding these error terms to the theoretical prices. Totally are 500 price series simulated yielding 500 pseudo distributions which are used to extract a 95 percent confidence band. The second method, bootstrap from actual error terms, is easier to implement. As in Söderlind [2000], the error terms are drawn with replacement from the current error terms, but to account for heteroscedastisity the error terms are grouped depending on option type and whether they are in the money or out of the money. The confidence bands are then obtained as described above.

Both methods produce fairly narrow confidence bands, but are less narrow than in the Monte Carlo experiments by Melick and Thomas [1998], Söderlind and Svensson [1997] and Söderlind [2000]. The confidence band obtained by bootstrapping from actual errors seems also to be wider than the counterpart of Söderlind [2000]. Andersson and Lomakka suggest that these differences are due to the non-normal and heteroscedastic features of the pricing errors. In the Monte Carlo experiments both these features are neglected, while in the bootstrap method of Söderlind [2000] only the heteroscedasticity is neglected. In line with the findings of Bliss and Panigirtzoglou [1999] and Cooper [1999] they conclude that the smoothed implied volatility smile method seems to be more robust than the double lognormal method.

As discussed, the confidence bands are primarily used to compare the robustness of different estimation techniques by studying the width of the bands. Andersson and Lomakka [2001] suggest an extended use of the confidence bands. In earlier studies, when assessing changes in market expectations due to specific economic events, conclusions have been drawn just by comparing the implied risk-neutral densities visually. No attempt have been made to quantify whether changes are statistically significant. Such a procedure is not satisfactory due to the noise attached to the estimated distributions. Andersson and Lomakka propose to use the confidence band to determine whether changes in the shape of the implied distributions are significant.

 $^{^{15}}relativestrike price = \frac{strike price}{future sprice} - 1$ 16 The smoothing procedure is performed in the implied volatility/delta space.

They classify an event as insignificant in statistical sense if the density estimated after the event falls within the 95 percent confidence band derived from the pre-event distribution. Similarly, if the density falls outside the confidence band, the opposite conclusion can be drawn.

The approach of Andersen and Lomakka for determining whether changes in the implied risk-neutral densities are significant or not is of great interest from a practitioner's point of view. For example, it might provide central banks with a more reliable indicator of whether market participants change their attitude towards risk when new information hits the market. In the present study, an analysis based on the methodology of Andersson and Lomakka [2001] is performed in the Norwegian option market. The aim is to gain a better understanding of how implied risk-neutral densities can be used to assess changes in market expectations in the Norwegian option market.

3.5 Information Content of Implied Risk-Neutral Density Functions

One of the main objects of this study is to evaluate whether properties of implied risk-neutral density functions extracted from OBX options can be used as leading indicators in the Norwegian equity market. Relatively few studies have examined the predictive capabilities of the information contained in implied distributions. A possible reason is that the density functions obtained from option prices are based on a risk-neutral pricing distribution which may differ significantly from the market participants' subjective density function. Thus, for forecasting purposes, the value of the information contained in risk-neutral distributions may be limited. A short review of selected studies is presented next.

Gemmill and Saflekos [2000] use a statistical measure called the hedge portfolio error to compare the forecasting performance of the double lognormal model relative to the Black-Scholes model. The hedge portfolio error is defined as the difference between the change of the market quoted option price and the change in the theoretical price implied by the model. The analysis is performed on FTSE-100 index options between 1987 and 1997. The results show that the the forecasting performance is improved when using the double lognormal model.

Navatte and Villa [2000] apply a quite different approach to test the information content of implied risk-neutral density functions. They study how standard deviation, skewness and kurtosis estimated from the implied distribution can be used to predict the corresponding realized sample moments in the remaining time to maturity. This is an extension of existing literature testing the forecasting ability of implied volatilites. The tests are performed by simply regressing the realized sample moments on the implied moments. The implied risk-neutral density functions are estimated by a Gram-Charlier series expansion of the normal distribution¹⁷ using long-term CAC 400 options. Navatte and Villa [2000] find that the two first moments, standard deviation and skewness, contain a substantial amount of information for future moments, while the implied kurtosis contains little information to predict future kurtosis.

The ability of the implied moments to predict the realized sample moments is also studied by Weinberg [2001]. In his study the implied risk-neutral density functions are estimated using the smoothed implied volatility smile method. Options on S&P 500 future, U.S.dollar/Japanese yen futures and U.S. dollar/deutche mark futures spanning the late 1980's through 1999 are employed in the analysis. Weinberg concludes that implied volatility predicts future realized volatility. However, the at-the-money Black-Scholes implied volatility performs slightly better in predicting future volatility than the implied volatility obtained using the smoothed implied volatility smile method. The ability of implied skewness to predict realized skewness is found to be poor. Tests of the predicting ability of implied kurtosis is not performed.

A less technical study on the information content of implied risk-neutral density functions is performed by Lomakka [2001] on OMX index options. He studies how two parameters, termed the *skewness-parameter* and the *uncertainty-parameter*, can be used as leading indicators in the Swedish stock market. The skewness parameter is defined as the probability of an increase in the OMX index of 10 percent or more, minus the probability of a decrease in the OMX index of 10 percent or more. The uncertainty parameter is defined as the sum of the above probabilities. The results show that the skewness-parameter is not able to predict the future development in the stock market. For the uncertainty-parameter, the result is slightly more positive. It seems to contain some information valuable for forecasting the future stock market development.

¹⁷Specify the normal distribution as a "prior" distribution and add correction terms to it by using a Gram-Charlier polynomial expansion.

Chapter 4

Methodology

In this chapter I will give a comprehensive description of the methodology used in the present study to extract and examine implied risk-neutral density functions. Two of the most popular methods for deriving implied distributions are applied, the double lognormal method (DLN) and the smoothed implied volatility smile method (SPLINE). In addition, the single lognormal model (SLN) is provided as a benchmark. I will first show how to implement these methods. I then present basic summary statistics used to analyse the distributions. The methodology applied for assessing the uncertainty of implied density functions extracted from OBX options is outlined next. Finally, I show how to test the predictive power of the information contained in OBX option prices.

4.1 The Single Lognormal Model (SLN)

In Chapter 2, the famous Black-Scholes model for valuing European call and put options is reviewed. This model assumes that the price of the underlying asset is lognormally distributed and that the return of the asset is normally distributed with constant variance. The lognormal distribution for a stochastic variable x can be described by two parameters, α and β , as:

$$L(x|\alpha,\beta) = \frac{1}{x\beta\sqrt{2\pi}}e^{-(\ln x - \alpha)^2/(2\beta^2)}$$

$$(4.1)$$

In the Black-Scholes model, options are priced as if investors are risk neutral by setting the expected rate of return on the underlying asset, μ , equal to the risk-free interest rate, r. The parameters of the risk-neutral lognormal distribution for the underlying asset at maturity can then be expressed as¹:

$$\alpha = \ln S_0 + (r - \frac{\sigma^2}{2})\tau \tag{4.2}$$

¹See Section 2.2.1, equation (2.7).
$$\beta = \sigma \sqrt{\tau} \tag{4.3}$$

where S_0 is the current price of the underlying asset, τ is the remaining time to expiration, and σ is the the volatility of the underlying asset.

Since the expected rate of return on the underlying asset is equal to the risk-free interest rate, the expected future value of the underlying asset at maturity must equal $S_0 e^{r\tau}$. The mean of the lognormal distribution is given by $e^{\alpha+\beta^2/2}$. Thus, the current price of the underlying asset, S_0 , can be found by²:

$$S_0 e^{r\tau} = \hat{E}[L(S_T | \alpha, \beta)] = e^{\alpha + \frac{1}{2}\beta^2}$$
 (4.4)

$$\Rightarrow S_0 = e^{-r\tau} e^{\alpha + \frac{1}{2}\beta^2} \tag{4.5}$$

Note that $\hat{E}()$ represents the expectation taken with respect to a risk-neutral distribution, which has expected return equal to the risk-free interest rate.

If we substitute the expression for S_0 and the expressions for α and β given in (4.2) and (4.3) into the Black-Scholes formulas given in equation (2.11) and (2.12), the price of an European call and put option with strike price X can be written as:

$$c(X,\tau) = e^{-r\tau} \left[e^{\alpha + \frac{1}{2}\beta^2} N(d_1) - XN(d_2) \right]$$
(4.6)

$$p(X,\tau) = e^{-r\tau} \left[-e^{\alpha + \frac{1}{2}\beta^2} N(-d_1) + XN(-d_2) \right]$$
(4.7)

where

$$d_1 = \frac{-\ln(X) + \alpha + \beta^2}{\beta} \tag{4.8}$$

$$d_2 = d_1 - \beta \tag{4.9}$$

Equations (4.6) and (4.7) are the general Black-Scholes formulas expressed in terms of the the parameters of the underlying lognormal distribution at maturity, α and β . If we assume a proxy for the risk free interest rate, these are the only unknown parameters. To extract the implied risk-neutral density function we need to estimate α and β . As described in Section 3.3.2, the distributional parameters can be estimated by minimizing the squared deviation between the observed option prices and the theoretical option prices calculated from (4.6) and (4.7). This is explained more in detail in the next section.

²Note that T is the time at maturity, and τ is the remaining time to maturity. In this case, $\tau = T - 0 = T$.

4.2 The Double Lognormal Model (DLN)

As pointed out in Chapter 2, options are generally not priced as though the price of the underlying asset is lognormally distributed. Hence, a more flexible density function is required. As described in Chapter 3, a widely used method is to assume that the distribution for the underlying asset is a weighted sum of several independent lognormal distributions. Since option contracts are generally traded for only a small range of strike prices, the number of parameters which can be extracted from observed option prices is limited. The number of available strikes is particularly low in the Norwegian market. To limit the number of estimation parameters I will assume that the underlying distribution is a combination of only *two* lognormal density functions. This distribution should be sufficiently flexible to capture characteristic features in the data such as fatness in the tails and positive or negative skewness.

The double lognormal distribution is described by five parameters: two parameters for each lognormal distribution (α_1 , β_1 and α_2 , β_2), and a weighting parameter (θ) which describes the relative weight of each distribution ($\theta \in \langle 0, 1 \rangle$). The risk-neutral double lognormal distribution can be written as:



$$q(S_T) = \theta \cdot L(S_T | \alpha_1, \beta_1) + (1 - \theta) \cdot L(S_T | \alpha_2, \beta_2)$$

$$(4.10)$$

Figure 4.1: Example of a double lognormal distribution

An illustration of a double lognormal distribution is given in Figure 4.1. The solid line is the double lognormal density which is equal to the sum of the two weighted lognormal components represented by the dashed and the dotted lines.

In the previous chapter, it is shown how to estimate the distributional parameters by minimizing the squared deviation between the observed option prices and the theoretical option prices obtained from the risk-neutral valuation equations. In practice, this minimizing problem is simplified. Bahra [1997] shows that if the price of the underlying asset is assumed to follow a double lognormal distribution, the price of an European call and put option can be expressed analytically. Hence, the minimizing problem can be simplified by using the analytical price expressions instead of the risk-neutral valuation equations. The expressions for the call and put price are:

$$c(X,\tau) = e^{-r\tau} \{\theta[e^{\alpha_1 + \frac{1}{2}\beta_1^2} N(d_1) - XN(d_2)] + (1-\theta)[e^{\alpha_2 + \frac{1}{2}\beta_2^2} N(d_3) - XN(d_4)]\}$$
(4.11)

$$p(X,\tau) = e^{-r\tau} \{\theta[-e^{\alpha_1 + \frac{1}{2}\beta_1^2}N(-d_1) + XN(-d_2)] + (1-\theta)[-e^{\alpha_2 + \frac{1}{2}\beta_2^2}N(-d_3) + XN(-d_4)]\}$$
(4.12)

where

$$d_1 = \frac{-ln(X) + \alpha_1 + \beta_1^2}{\beta_1} \quad , \ d_2 = d_1 - \beta_1 \tag{4.13}$$

$$d_3 = \frac{-\ln(X) + \alpha_2 + \beta_2^2}{\beta_2} \quad , \ d_4 = d_3 - \beta_2 \tag{4.14}$$

If we compare the option price formulas in (4.11) and (4.12) with the expressions obtained for the single lognormal model in (4.6) and (4.7), we see that under the double lognormal assumption the option price is a weighted sum of two Black-Scholes solutions.

In the absence of arbitrage opportunities, the currently observed futures price (F) for the underlying asset should equal the mean of the risk-neutral density function. The mean of the double lognormal distribution is a weighted sum of the individuals means. Hence, the following relationship must be satisfied:

$$\theta e^{\alpha_1 + \frac{1}{2}\beta_1^2} + (1 - \theta)e^{\alpha_2 + \frac{1}{2}\beta_2^2} = F \tag{4.15}$$

One way of implementing (4.15) is to impose it as a constraint in the minimizing problem. However, this constraint may be binding and thus reduce the goodness-of-fit. A more general method is to add the squared deviation between the futures price and the mean of the distribution to the minimization problem.

If we assume a proxy for the risk free interest rate, we have five distributional parameters that need to be estimated in order to obtain the implied risk-neutral density function. These parameters are recovered by minimizing the squared deviation between the observed option prices and the theoretical option prices calculated from (4.11) and (4.12) across all exercise prices for a given maturity, together with the squared mean-futures price deviation. The same procedure is applied for the single lognormal model. The minimization problem can be written as:

$$\min_{\alpha_1,\alpha_2,\beta_1,\beta_2,\theta} \{ \sum_{i=1}^m (c_i - c_i^*)^2 + \sum_{j=1}^n (p_j - p_j^*)^2 + [\theta e^{\alpha_1 + \frac{1}{2}\beta_1^2} + (1 - \theta)e^{\alpha_2 + \frac{1}{2}\beta_2^2} - F]^2 \}$$
(4.16)

where

 $c_i^*, p_j^* =$ Observed option prices $c_i, p_j =$ Theoretical option prices estimated from (4.11) and (4.12) m =Number of call options in the data set n =Number of put options in the data set

As pointed out in Section 2.2.4, European options on futures and on spot with the same strike price and time to maturity are in theory equivalent. This is because the underlying variable is equal at maturity ($F_T = S_T$). Consequently, the underlying risk-neutral distribution is also equal at maturity. Thus, when recovering the distributional parameters, using futures options or spot options are equivalent.

The minimization problem in (4.16) is implemented in MATLAB. First, a MATLAB-routine is written which calculates the sum of the squared price deviations for given values of α_1 , α_2 , β_1 , β_2 and θ . Then, a built-in optimization procedure in MATLAB is employed to find the optimal parameters that yield the smallest sum of squared deviations.

The double lognormal method has some weaknesses. It sometimes produces a density function which is characterized by a sharp spike. The reason is that one of the two lognormal distributions is estimated to have a very small standard deviation. It also happens that the optimization procedure fails in finding any solution for a particular dataset. To overcome these problems I have put the following restrictions on the β values ³:

$$0.25 < \frac{\beta_1}{\beta_2} < 4 \tag{4.17}$$

4.3 The Smoothed Implied Volatility Smile Method (SPLINE)

If we are able to express the call price function as continuous function of the strike price, the risk-neutral density function can be obtained easily using the Breeden and Litzenberger [1978] result⁴. In Section 3.3.3 it is explained how Shimko [1993] derives the call price function by first interpolating the volatility smile, and then substituting the expression for the volatility in the Black-Scholes model. The same methodology, with minor modifications, is used in the present study. I follow Malz [1997] in interpolating the implied volatilities across deltas in stead of

 $^{^3 \}rm Suggested$ by Andersson and Lomakka [2001].

 $[\]frac{4 \partial^2 C(X,\tau)}{\partial X^2} = e^{-r\tau} q(S_T)$ (see Section 3.2).

across strike prices, and Campa et al. [1998] in using smoothing splines to fit the function⁵. The general procedure is outlined first, then the various components of the method are considered in more detail.

4.3.1 General Procedure

In the smoothed implied volatility smile method, we start out by calculating the implied volatilities and the delta values from the observed set of option prices and strike prices. A simple iterative procedure written in MATLAB is used to back out the implied volatilities from the Black-76 model. The Black-76 model is used because index options generally are priced as though they are options on the index future⁶. The corresponding deltas are then estimated according to:

$$\Delta = \frac{\partial c}{\partial F} = e^{-r\tau} N(d_1), \quad d_1 = \frac{\ln(\frac{F_0}{X}) + \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}}$$
(4.18)

The implied volatilities are plotted against delta values, and a smoothing spline function is fitted to the data. We are now able to express the volatility as a continuous function of delta, $\sigma(\Delta)$. The next step is to transform the delta values to their corresponding strike prices by using the definition for d_1 :

$$X = \frac{F_0}{e^{(\sigma(\Delta)\sqrt{\tau}N^{-1}(e^{r\tau}\Delta) - \frac{1}{2}\sigma(\Delta)^2\tau)}}$$
(4.19)

Now we can express the volatility as a continuous function of the strike price, $\sigma(X)$, and the call pricing function, c(X), is obtained by substituting this expression in the Black-76 model. The implied risk-neutral density function is then derived by differentiating the call pricing function twice with respect to the strike price and multiplying it by the inverse discount factor. The differentiation is done numerically using a simple finite difference approximation:

$$\frac{\partial^2 c}{\partial X^2} = \frac{c_{i+1} + c_{i-1} - 2c_i}{(\triangle X)^2}$$
(4.20)

The finite difference approximation in (4.20) requires constant intervals between the strike prices. To achieve this, equation (4.19) is used to iterate out delta values corresponding to a set of equally spaced strike prices. The smoothed implied volatility smile method is illustrated in Figure 4.2.

4.3.2 The Smoothing Spline Procedure

We have seen that in order to establish a complete risk-neutral distribution we need to connect the discrete observations into a continuous function. One way of connecting the observations

⁵The rationale for these modifications is explained i Section 3.3.3.

⁶See discussion in Section 2.2.4.



Figure 4.2: Steps in the smoothed implied volatility smile method

together is to fit a single polynomial to the discrete observations. This is a very simple method, and may not fit the data very well. Instead of using a single polynomial on the entire interval, we can obtain a more precise and less restrictive function by fitting several polynomials to the observations. This technique is called splines.

In the present study, cubic splines are used to interpolate the data. Cubic splines are piecewise cubic polynomials defined such that the function and its first two derivatives are continuous at the knot points. Since we are working with noisy data, an exact interpolation of the observed points may give a highly oscillating curve. Such excessive oscillations are not "reasonable" for the implied volatiliy curve. To reduce the oscillations and increase the smoothness of the cubic spline, we apply so called smoothing splines. Smoothing splines damps out the oscillations by seeking the solution to the modified sum of squares expressed as⁷:

$$\min_{\Phi} [\lambda \sum_{i} w_i [y_i - f(x_i; \Phi)]^2 + (1 - \lambda) \int f''(x; \Phi)^2 dx]$$
(4.21)

where x_i and y_i are the observations to be interpolated (i.e. delta values and implied volatilities, respectively), $f(x; \Phi)$ is the spline function, Φ is the parameter matrix of the spline function, w_i is the weighting parameter for observation *i*, and λ is the smoothing parameter. The smoothing spline procedure in (4.21) is implemented using the Spline Toolbox in MATLAB.

The minimizing problem in (4.21) consists of two parts. The first part minimizes the squared deviations between the observed values and the values generated by the spline function. This part controls the goodness-of-fit of the spline function. The second part minimizes the integral on the squared curvature of the spline function. As the variability of the function increase, this part will increase too. Thus, the second part controls the smoothness of the spline function. The smoothing parameter, λ , decide the relative weight of the two parts, and thus controls the trade off between minimizing the residual error and minimizing local variation.

The smoothing parameter is of great importance in the smoothing spline procedure. A large value of the smoothing parameter means that the minimization procedure puts greater weight on minimizing the residual errors. Equivalently, for a low value, greater weight is put on minimizing the curvature. A smoothing parameter equal to one $(\lambda = 1)$ means that the spline function would accurately interpolate the data. On the other extreme, a smoothing parameter equal to zero $(\lambda = 0)$, gives a function which minimizes the curvature. In this case the spline function becomes a straight line (i.e. a linear regression).

Choosing the "Optimal" Smoothing Parameter

Choosing the appropriate smoothness parameter is an important step in practice. Choosing a too low value may give a spline function that is too smooth and does not fit the data very well. A such function may ignore important characteristics of the data. On the other hand, a too high value may give a resulting density function with too much oscillation. Even though the spline function in the implied volatility/delta space seems smooth, the transformation from $\sigma(\Delta)$ to c(X), and then the twice-differentiation to obtain the implied risk-neutral density may give a distribution that is not particularly smooth. The role of the smoothing parameter is illustrated in Figure 4.3. The solid line represents a smooth density function. The dashed and the dotted lines are density functions obtained by using a larger smoothing parameter. Wee see that as we increase the smoothing parameter, the smoothing is reduced and the resulting density function becomes more oscillating.

⁷See for example Bliss and Panigirtzoglou [1999], Appendix A.2.



Figure 4.3: An implied risk-neutral density function for various smoothing parameters

A suitable smoothing parameter can be obtained by simply plotting the distribution for different smoothing parameters and choose the one which yields the "best" result. This is of course a cumbersome method if many distributions need to be estimated. Several methods have been proposed to determine the "optimal" smoothing parameter. One well-known method is the cross-validation score method proposed by Craven and Wahba [1979]⁸. The idea behind this method is to exclude the observations one by one, and find the smoothing parameter value under which the missing data points are best predicted by the remainder of the data. More precisely, for a given smoothing parameter the observations are deleted one by one and a spline function is estimated in each case from the remaining observations. The sum of the squared errors between the deleted observations and the values generated by the spline functions is then calculated. This procedure is repeated for all possible smoothing parameters. The "optimal" smoothing parameter is the one that yields the smallest sum of squared errors. The cross-validation method can be written as:

$$\min CVS(\lambda) = \sum_{i} (y_i - g_{\lambda[-i]}(x_i))^2$$
(4.22)

where x_i and y_i are the actual observations and $g_{\lambda[-i]}(x_i)$ is the smoothing spline function estimated by excluding observation (x_i, y_i) .

For the dataset applied in this study, the cross-validation algorithm sometimes computes values of λ that result in either oscillating risk-neutral density functions or distributions that are too smooth and do not fully capture the characteristics of the data. This indicates that even if the spline function is optimal according to the cross-validation procedure in the implied volatility/delta space, it is not necessarily "optimal" after the transformation required to obtain the risk-neutral density function. Due to this lack of consistency in the cross-validation procedure, the smoothing parameter is chosen manually by picking the highest smoothing parameter that

⁸Applied for example by Andersson and Lomakka [2001].

produces a non-oscillating risk neutral density function. Generally, the smoothing parameter is set around 0.998^9 .

Weighting

The smoothing spline procedure allows us to set the relative weights of each observation. Bliss and Panigirtzoglou [1999] discuss different types of weighting schemes and how the weighting can account for different sources of pricing error. In the smoothing spline method they propose to weight each parameter in terms of the Black-Scholes parameter vega, v^{10} . This implies that most weight are placed on options near at the money, and less weight on options away from at the money. The v-weighting scheme is compared with other alternative types of weighting, and the results show that the weights have little impact on the estimation. Bliss and Panigirtzoglou [1999] explain this by the fact that the fitted price errors are generally very small.

In the present analysis I have chosen to use equal weighting of the observations in the smoothed implied volatility smile method. This is also applied in the double lognormal method and the single lognormal model. Since any weighting schemes assume a special structure for the option pricing errors, the weighting will be more or less ad-hoc. A such ad hoc weighting may easily introduce more noise in the estimation.

4.3.3 The Tails of the Distribution

As emphasized previously, option prices can only provide information about the underlying density at their respective strike prices. Generally, the range of available strike prices is limited to an area around at the money. Thus, a central question is how to extrapolate beyond the range of traded strike prices.

When using the implied volatility/delta space instead of the implied volatility/strike space to interpolate across the volatility smile, the options close to being at the money are less grouped together than options far away from at the money. Thus, a greater "shape" is permitted near the centre of the distribution. In addition, since possible values in the delta space range from 0 to $e^{-r\tau}$, the extrapolation area is generally relatively small. If we look back at Figure 4.2, it illustrates the difference between the implied volatility/delta space and the implied volatility/strike price space.

In the present study I test three different approaches for modeling the tails of the implied distribution:

⁹The smoothness of the density functions appeared to be very sensitive to small changes in the smoothing parameter above 0.998.

 $^{{}^{10}}v = \frac{\partial c}{\partial \sigma}$ is defined as the derivative of the Black-Scholes call price with respect to the Black-Scholes implied volatility.

- Use the first and last polynomial of the spline function to extend it left and right, respectively.
- Assume that the spline function is linear outside the range of observation.
- Fit lognormal distributions at the tails of the implied density function.

The first approach is motivated by the fact that this is the simplest method if using the Spline Toolbox in MATLAB. The spline function is here automatically extended left and right by using the last polynomial in each end. The second approach is in line with Bliss and Panigirtzoglou [1999]. Fitting lognormal distributions at the tails of the implied density function is equivalent to assuming that the volatility smile is flat outside the range of observations. This was first proposed by Shimko [1993], and is among others applied by Andersson and Lomakka [2001]. In the literature studied, the procedure of how to implement this last approach is not explicitly explained. Thus, the procedure applied here is suggested by the author.

A general problem when assuming that the volatility is constant outside the range of observations is that a smooth transition to the extrapolation area is needed. A too abrupt change in the implied volatility curve may cause problems with the differentiation of the call price function, and give a discontinuous density function. An illustration is given in Figure 4.4. In this example, the implied volatility curve is extended horizontally at each ends without any form of smoothing. We see that the break in the implied volatility curve leads to a discontinuity in the density function at the transition to the extrapolation area.



Figure 4.4: Implications of a too abrupt change in the implied volatility curve

To ensure a smooth transition, I have chosen to do this directly on the implied density function. The approach is illustrated in Figure 4.5. First, the implied distribution is estimated for the region of observable strike prices. Then, the left and the right tail of the distribution are estimated assuming a constant volatility equal to the (smoothed) volatility at the respective end observation. The three parts are then connected together using splines in the transition areas (dotted lines). The distribution is then scaled to ensure that the implied distribution integrates up to one¹¹.



Figure 4.5: Connecting lognormal distributions at the tails of the implied risk-neutral density function

4.4 Summary Statistics for Implied Risk-Neutral Density Functions

The implied risk-neutral density functions can be described by computing a range of summary statistics. These measures are useful when analyzing changes in the shape of the implied distributions, and when comparing different estimation techniques. The summary statistics applied in the present study are presented next.

The expected (risk-neutral) future value of the underlying asset is equal to the *mean* of the implied distribution. The mean is also sometimes referred to as the first *moment* of the distribution. As emphasized earlier, the mean of the risk-neutral density function theoretically equals the futures price. The mean of a density function f(x) is calculated as:

$$\mu = \int_{-\infty}^{\infty} x f(x) dx \tag{4.23}$$

Two other useful point measures are the *mode* and the *median*. Mode is the most likely future outcome of the distribution, while the median is the middle value of the distribution, i.e. half the probability mass is above the median and half is below.

¹¹Scaling is also necessary for the other two approaches.

The second moment of the distribution is the *variance*. The square root of the variance is the *standard deviation*. Standard deviation is a measure of the dispersion of the distribution around the mean and is calculated from:

$$\sigma = \sqrt{\int_{-\infty}^{\infty} (x-\mu)^2 f(x) dx}$$
(4.24)

A high standard deviation for the implied distribution indicates that there is a great uncertainty among the market participants about how the underlying asset will evolve towards maturity. Throughout this study the standard deviation is calculated from the return distribution and not the price distribution. The reason is that the standard deviation of the price might vary with the price level, and thus invalidates comparing the standard deviations at various dates if the price level is different. The standard deviation of return is on the other hand not generally depending on the price level, which makes it possible to compare the values at different dates.

Another measure describing the shape of the distribution is *skewness*. Skewness is calculated as the normalized third moment:

$$Sk = \frac{1}{\sigma^3} \int_{-\infty}^{\infty} (x-\mu)^3 f(x) dx \qquad (4.25)$$

Skewness characterizes the asymmetry of a density function. A normal distribution has zero skewness. If the implied distribution is positively skewed, the right tail is greater than the left tail. This may suggest that the market participants are positive about the future development. However, it is important to keep in mind that such positive expectations naturally lead to an upward revision of the mean/futures price (and consequently the stock price). Thus, in a positively skewed distribution there is less probability attached to outcomes higher than the mean than to outcomes below the mean, i.e. the median is lower than the mean.

Skewness is very sensitive to the tails of the distribution. As emphasized previously, the limited range of available strike prices and the generally reduced liquidity for strike prices away from the current futures price, introduce a high level of uncertainty in the tails of the implied risk-neutral density functions. This decreases the reliability of the standard measure of skewness. I have therefore also applied a skewness measure that is less sensitive to the tails of the distributions, the Pearson median skewness:

$$Pearson\ median\ skewness\ =\ \frac{\mu - median}{\sigma} \tag{4.26}$$

A measure of how peaked a distribution is and the likelihood of extreme outcomes is *kurtosis*. This is the normalized forth moment of the density function:

$$K = \frac{1}{\sigma^4} \int_{-\infty}^{\infty} (x - \mu)^4 f(x) dx$$
 (4.27)

The kurtosis of the normal distribution is equal to 3. It is therefore usual to calculate the *excess kurtosis*, which is defined to be the kurtosis less 3. A positive excess kurtosis implies a greater probability for extreme changes compared to the normal distribution. This means that the distribution has fatter tails than the normal distribution. Due to the uncertainty related to the shape of the tails of the implied risk-neutral density functions, the reliability of the kurtosis measure is poor and should be interpreted with care.

In addition to calculating the various moments of the distributions, I have also applied two probability measures which I have called *skewness-parameter* and *uncertainty-parameter*¹². The skewness-parameter is defined as the probability of the OBX index to exceed the futures price by 7.5 percent or more, minus the probability of the OBX index to fall below the futures price by 7.5 percent or more¹³. The uncertainty parameter is defined as the sum of the above probabilities. The skewness-parameter is a measure of the asymmetry of the distribution, while the uncertainty-parameter is a measure of the dispersion of the distribution, in probabilistic terms. These measures may intuitively be easier to interpret than skewness and kurtosis.

Calculations of the summary statistics are performed numerically using a simple trapezoidal integration rule¹⁴. The function to be integrated, f, is then approximated by n trapezoids. The approximation of the integral is calculated as:

$$J = \int_{a}^{b} f(x)dx \approx h\left[\frac{1}{2}f(a) + f(x_{1}) + f(x_{2}) + \dots + f(x_{n-1}) + \frac{1}{2}f(b)\right]$$
(4.28)

where h = (b - a)/n. In the present calculations, n is roughly 700.

An important question when analyzing the implied risk-neutral density function is how good the various models fits the option data they are supposed to describe. As a measure of the in-sample goodness-of-fit, I have calculated the root mean squared error (RMSE) between observed option prices and implied option prices from the models. The root mean squared error is calculated as:

$$RMSE = \sqrt{1/(m+n) \left[\sum_{i=1}^{m} (c_i - c_i^*)^2 + \sum_{j=1}^{n} (p_j - p_j^*)^2 \right]}$$
(4.29)

¹²Based on measures applied in Lomakka [2001].

 $^{^{13}}$ Using 7.5 percent, instead of for example 10 percent, is done to avoid ending up in the extreme tails of the distributions.

¹⁴See for example Kreyszig [1993], Chapter 18.

where

c_i^*, p_j^*	=	Observed option prices
c_i, p_j	=	Theoretical option prices estimated from implied risk-neutral density function
m	=	Number of call options in the data set
n	=	Number of put options in the data set

The implied option prices are calculated using the analytical expressions in (4.11) and (4.12) for the double lognormal model, and (4.6) and (4.7) for the single lognormal model. For the SPLINE method, the implied option prices are estimated numerically using the risk-neutral valuation equations in (3.2) and (3.3).

4.5 Bid-Ask Bounds

A major weakness of the option contracts traded at Oslo Stock Exchange is their low liquidity. Relatively large bid-ask spreads are observed in the data. In this study I am using an average of the closing bid and ask quotation as a proxy for the option price. Thus, a main source of error in the risk-neutral density estimation is the high level of uncertainty in input prices caused by the bid-ask spreads¹⁵. It is therefore of great interest to study the effect of the bid-ask spreads on the estimated distributions.

As reviewed in Section 3.4.2, the literature focusing on the uncertainty of the estimated riskneutral densities is perturbing either the observed or the theoretical prices in some appropriate manner, and analyses how this affects the estimation of the implied distributions. For example, Bliss and Panigirtzoglou [1999] perturb the observed option prices by a random number uniformly distributed from minus and plus one half of the contracts "tick size", while Andersson and Lomakka [2001] use historical or actual price errors to perturb the theoretical prices. In line with these studies, I will assess the uncertainty related to the bid-ask spreads by perturbing the option prices using the differences between the bid and the ask quotations. The object is also to study the robustness of the different estimation methods to disturbances in option prices.

The simulated price series are simply created by drawing uniformly from the bid-ask spreads. More precisely, for each strike price I obtain a new option price by drawing randomly from a uniform distribution between the bid and the ask price. This is repeated 500 times. From the simulated option price series, I calculate 500 distributions and estimate the summary statistics for each of them. I can then study the dispersion of the summary statistics by estimating the 5^{th} percentile and the 95^{th} percentile for the samples of the various summary statistics. To visually illustrate the effect of the bid-ask spreads, I derive confidence bands based on the 500 simulated distributions. A confidence band is derived by calculating the 5^{th} percentile and the

 $^{^{15}}$ See discussion in Section 3.4.1.

 95^{th} percentile for the density at each value along the horizontal axis. Thus, the confidence band defines the inner and outer bound for each density point at a 90 percent significance level. Notice that the inner and outer bounds do not represent possible density functions. Since the error band represents the confidence intervals for each single density point, the lower bound necessarily integrates to less than unity, and the upper bound integrates to more than unity.

4.6 Testing the Information Content of Option Prices

The recent interest in developing techniques for extracting information about market expectations from option prices is related to the forward looking nature of option prices. Thus, an important question is whether properties of implied risk-neutral density functions can be used as leading indicators of the future development in the financial markets.

As reviewed in Section 3.5, earlier literature studying the predictive power of information contained in option prices has been focusing on the relationship between the implied moments and the corresponding realized sample moments¹⁶. For example, can volatility and skewness estimated from the implied risk-neutral density functions be used to predict realized future volatility and skewness? For implied volatility, an approximate linkage exist between estimates based on daily returns and estimates based on the risk-neutral density function. But for skewness, and higher moments, the linkage is more complex¹⁷. Thus, in this study I will only consider whether implied volatility can be used to predict future volatility. Similar tests for implied skewness is not performed. To evaluate the information content of the asymmetry of the implied distributions, I will study how properties of the implied risk-neutral density function can be used to predict the deviation between the realized spot price at maturity and the futures price at the observation date. The analysis is based on a simple non-parametric sign test.

4.6.1 Realized vs Implied Volatility

The first test to be performed is whether the standard deviation estimated from the implied riskneutral density function can be used to predict the realized standard deviation of the underlying asset in the remaining time to maturity. The methodology is based on Weinberg [2001].

As the time to maturity increases, the uncertainty about how the price of the underlying asset will evolve towards maturity will naturally also increase. Thus, the standard deviation of the risk-neutral density function is positively related to the remaining time to maturity. To be able to compare the standard deviation estimated from an implied distribution with a given time to maturity, and the realized volatility based on daily returns, the standard deviation from the implied risk-neutral density function is re-calculated as:

¹⁶See for example Navatte and Villa [2000] and Weinberg [2001].

¹⁷See Weinberg [2001].

$$\sigma^* = \sqrt{\ln(\frac{\sigma^2}{\mu^2} + 1)/\tau} \tag{4.30}$$

Equation (4.30) was first presented by Jarrow and Rudd [1982]. The underlying assumption is that the second moment (but not the higher moments) of the risk-neutral density function can be approximately represented by the lognormal property of the Black-Scholes model. Note that σ and μ are the standard deviation and the mean, respectively, calculated from the implied *price* distribution, and not the return distribution. τ is the remaining time to maturity. The realized future volatility is estimated over the remaining trading days of the option contract, τ^* , by:

$$\sigma_f = \sqrt{1/\tau^* \sum_{i=1}^{\tau^*+1} (R_i - \bar{R})^2}$$
(4.31)

 R_i is the logarithmic return on day i and \overline{R} is the mean return over the remaining days of the contract.

To test the ability of implied volatility to predict realized volatility, a simple ordinary least squares (OLS) regression is applied:

$$\sigma_{f,t} = a + b \cdot \sigma_t^* + e_t \tag{4.32}$$

If the volatility calculated from the implied risk-neutral density functions contains some useful predictive information about future volatility, b must be significant and positive. a = 0 and b = 1 corresponds to the hypothesis that implied volatility is an unbiased forecast of future volatility.

4.6.2 Information Content of Implied Skewness

To test whether the asymmetry, or the skewness, of the implied risk-neutral density functions can provide predictive information, I propose a simple non-parametric sign test based on the deviation between the spot price at maturity and the futures price at the observation date.

Futures Prices vs Expected Future Spot Price

An often raised question in the finance literature is whether the futures price of an asset equals its expected future spot price. Early literature on this subject focus on the relationship between the short and long positions of hedgers and speculators to explain deviations between futures prices and expected future spot prices¹⁸. If hedgers tend to hold short positions (selling futures), and speculators tend to hold long positions (buying futures), it is argued that the futures price

¹⁸See discussion in Hull [2000] Section 3.12.

will be below the expected future spot price since speculators require compensation for bearing the risk of the hegders. When the futures price is below the expected future spot price, it is known as *normal backwardation*. On the other hand, if speculators tend to hold short positions and hedgers long positions, the futures price must be above the expected future spot price. This is known as *contango*. Only when there is balance between short hegders and long hedgers in the market, the futures price will be equal to the expected future spot price.

Normal bacwardation and contango can also be explained by considering the tradeoffs between risk and return in capital markets. Under given assumptions, it can be shown that if the price of the underlying asset is uncorrelated with the level of the stock market, the futures price will equal the expected future spot price. Normal backwardation and contango occurs if the price is positively and negatively correlated with the level of the stock market, respectively.

A lot of empirical work have been carried out testing the various hypothesizes concerning the relationship between the futures price and the expected future spot price. The conclusions from these studies are mixed¹⁹.

Testing the Predictive Power of Skewness

As a starting point for the analysis, I assume that the futures price is the market participants best guess on the future spot price at maturity, and that the implied risk-neutral density function contains valuable information about how the market participants assess the uncertainty attached to the point estimate of the future spot price. The hypothesis to be tested is whether the deviation between the futures price observed today and the spot price observed at maturity is related to the asymmetry of the implied risk-neutral density function.

If the implied risk-neutral density function for example exhibits a pronounced right tail, the probability mass above the expected value is smaller than the probability mass below the expected value. This is illustrated in Figure 4.6. In other words, since the mean of the riskneutral density function is equal to the futures price, a positively skewed distribution implies that there is a greater probability for the future spot price to be below the futures price observed today than above. For a negatively skewed distribution, the situation is reversed.

To test whether the deviation between the futures price observed today and the spot price observed at maturity can be predicted using properties of the implied risk-neutral density function, I will apply a simple non-parametric *sign test*. The test parameter, z, is calculated as:

$$z = \begin{cases} 1 & \text{if } Q_t > 0 \text{ and } S_{T,t} - F_{0,t} > 0 \\ & \text{or } Q_t < 0 \text{ and } S_{T,t} - F_{0,t} < 0 \\ 0 & \text{otherwise} \end{cases}$$

¹⁹See discussion in Hull [2000] Section 3.12.



Figure 4.6: A positively skewed distribution

where $S_{T,t}$ is the spot price at maturity for contract t and $F_{0,t}$ is the corresponding observed futures price at the date when the implied risk-neutral density function is estimated. Q is a parameter derived from the implied distribution which is assumed to explain the deviation between the futures price and the future spot price (defined below). When Q is positive, the future spot price is expected to be above the futures price, and vice versa. Thus, z = 1 if the prediction of the model is correct, and z = 0 if it is wrong. By counting the number of ones we know how many times the prediction of the model is correct. Under the null hypothesis of no predictive power, the number of ones should not be significantly greater than number of zeros. If the model has any predictive power, the number of ones should be significantly higher than number of zeros. An evident weakness of this non-parametric sign test is of course that it gives us no information about the strength of the relationship we are testing, only whether a relationship exist or not.

Three different definitions of Q are applied in the analysis. One is based on standard skewness, another on Pearson median skewness, while the third is based on the skewness-parameter. These measures captures the asymmetry of the implied risk-neutral density function, and may contain valuable information about the market participants' assessment of the direction in the future development. Notice an important distinction between the skewness-parameter, and the other two skewness measures. Skewness and Pearson median skewness consider the relative size of the *total* probability mass at each side of the mean/futures price. The skewness-parameter, on the other hand, measures the excess probability of a *large* positive price movement relative to a *large* negative price movement. Thus, it only considers *parts* of the probability mass at each side of the futures price. Notice also that the signs of skewness and Pearson median skewness have to be reversed when using these measures in the definition of Q. When these measures are positive, the implied distribution is positively skewed, which means that a greater part of the probability mass is located below the mean/futures price than above. This implies that the likelihood of the future spot price to be below the futures price is greater than the likelihood of being above.

The measures are scaled relative to a "normal" level such that when they exceed the "normal" level (Q > 0), the future spot price is expected to rise above the futures price, and when they fall below the "normal" level (Q < 0), the future spot price is expected to be less than the futures price. An important question is what to define as the "normal" level. Obvious candidates are for example the mean or the median of the variable. An alternative approach is to define the normal level for each measure such that the number of Q > 0 match the number of $(S_T - F_0) > 0$, and vice versa. This means that number of positive and negative values for Q and $(S_T - F_0)$ are equally distributed. The latter approach will clearly lead to a stronger relationship between the futures price and the future spot price. But for testing purposes, it may be a good starting point. If this approach does not work, nothing works.

Chapter 5

Analyses

The object of the present analyses is to study the behaviour of implied risk-neutral density functions in the Norwegian option market. I will first give examples of implied distributions during a financial stress event. The purpose is to illustrate how the distributions may change during such events. I then compare how the double lognormal method and the smoothed implied volatility smile method perform relative to each other, and relative to the standard lognormal model underlying the Black-Scholes model. Another important question that will be considered is how the bid-ask spreads affect the estimation of the implied distributions. Finally, I am going to study whether properties of implied risk-neutral density functions can be used as leading indicators in the Norwegian stock market.

5.1 Data

The study is performed on equity options on the OBX index at Oslo Stock Exchange from December 1995 until January 2002. The OBX index consists of the 25 most traded stocks on Oslo Stock Exchange the last 6 months, and may therefore be a good indicator for the total Norwegian stock market. The options are European style. As mentioned in Section 2.2.4, futures contracts are also traded on the OBX index. Since the futures contract matures at the same time as the option contract, the futures can be used as a proxy for the underlying OBX index.

The option contracts at the OBX index are written with three months to expiration. A new contract is introduced every month. This means that there are always three different option contracts traded at the OBX index. Unfortunately, the liquidity in the Norwegian option market is rather poor. It is therefore recommended to use the next-to-expiration contract when extracting risk-neutral density functions. Thus in this study, options with four weeks to maturity are employed in the estimation.

The OBX option data provided by Oslo Stock Exchange consist of daily bid and ask quotations at closing time and closing prices. As discussed in Section 3.4.1, using closing prices may lead to problems with non-synchronicity. Closing prices at the end of the day are likely to be non-synchronous due to the infrequent trade of most options and the great variations in time-of-last-trade. In addition, the sample may include strikes that have not been traded during the trading day. The problem of non-synchronicity is naturally a great concern in the Norwegian option market due to the poor liquidity. Thus, using closing prices in the density estimation may give spurious results. To obtain a more reliable proxy for the option price, I have used an average of the closing bid and ask quotation. Relatively large bid-ask spreads are observed in the data, so we have to keep in mind that using an average of the bid and ask quotation is a rough approximation of the underlying economic price, especially for deep out-of-the-money options were the bid-ask spreads become a higher percentage of the option premium. An average of the closing bid and ask quotation for the OBX futures is used as a proxy for the futures price. This approximation is of minor concern due to the small bid-ask spreads observed for the futures price.

The option price data are filtered to eliminate observations that allow for obvious arbitrage opportunities. The filtering ensures that the price of a call option does not become more expensive for higher strike prices, and that a put option does not become less expensive for higher strike prices. Observations that produce an implied volatility of zero is also eliminated from the data set. The number of available option prices retained for analysis on each observation date (calls and puts) varies between six and twenty-three.

Both put and call options may be available for a given strike price. Since put-call parity¹ does not generally hold for the data set analysed in the present study, these contracts are not redundant. In the smoothed implied volatility smile method only one volatility per strike price can be utilized. The conventional wisdom in the major derivative markets is that an out-of-the-money option is more liquid that its in-the-money counterpart². So if both a put and a call are traded for a given strike price, the out-of-the-money option is selected. In the Norwegian option market, call options are generally more liquid than put options. So one might argue that the call option should be selected when both a put option and a call option are available for a given strike price. But to avoid any elimination of data, I have chosen to use the average of the call and put volatility when both options exist for a given strike price. The number of observations retained for the smoothed implied volatility smile method on each date varies between five and thirteen.

As a proxy for the risk-free interest rate I have used the 1 month Norwegian currency swap rate, which is assumed to be the most liquid interest rate product in the Norwegian market. The swap rate is adjusted for credit risk. The credit risk is assumed to be 20 basis points³. However,

¹The theoretical relationship between the price of an European put and an European call on the same underlying asset with the same expiration date, which prevents arbitrage opportunities $(c + Xe^{-r\tau} = p + S_0)$.

²See for example Bliss and Panigirtzoglou [1999].

³Based on discussions with people in the Market Operation Department in Norges Bank.

small variations in the interest rate have a negligible effect on option prices⁴.

5.2 Examples of Implied Risk-Neutral Density Functions in the Norwegian Option Market

In this section I will give examples of implied risk-neutral density functions in the Norwegian option market. The distributions are extracted around a financial stress event to illustrate how the shape of distributions may change in response to such events. The examples also illustrate how the implied risk-neutral density functions can be used to extract information about changes in market expectations.

The development in the OBX index from January 1994 to February 2001 is illustrated in Figure 5.1. In this period there are three main stress events: the Asian crisis in 1997, the Russian crisis in 1998, and the collapse in the IT-sector in 2001 followed by the attack on World Trade Center at September 11. As an example, I will focus on the Russian crisis. Options with four weeks to maturity are employed to extract the risk-neutral density functions. This means that the implied distributions reflect the market participants' expectations about the future development in the OBX index four weeks ahead. It is important to realize that when comparing two distributions, the remaining time to maturity should be the same. This is because the uncertainty about how the price of the underlying asset will evolve towards maturity is reduced as the remaining time to maturity decreases.



Figure 5.1: The OBX index from January 1994 to February 2002

The implied risk-neutral density functions are extracted using the double lognormal method (DLN) and the smoothed implied volatility smile method⁵ (SPLINE). In addition, a standard

⁴See for example Rubinstein [1985].

⁵Using a linear extrapolation of the spline function outside the observed range of delta values.

single lognormal model is presented as a benchmark (SLN). The distributions are presented with logarithmic return on the horizontal axis, or more precisely, logarithmic return relative to the future price, i.e. $ln(S_T/F_0)$. This means that the return distribution produced by the single lognormal model is a normal distribution. To simplify the interpretation of the implied distributions, I have calculated standard deviation, skewness and implied risk-neutral probability of a fall in the OBX index relative to the futures price of more than 7.5 percent for each distribution. Note that the skewness of the implied (logarithmic) return distribution obtained by the single lognormal model is zero. As a measure of how good the various models fit the option data, I have calculated the root mean squared error (RMSE) between observed option prices and implied option prices.

5.2.1 The Russian Crisis

In the second half of 1997 the Norwegian stock market experienced a significant fall as a result of the financial crisis that took place in several Asian countries. However, the Asian crisis led only to a minor correction in the Norwegian market, and it soon recovered to a new high level. But there was more to follow. The turbulence experienced during the financial crisis in Russia in 1998 was far more dramatic. From April 28 to October 8, a period of less than six months, the OBX index fell from 769 to 395, an almost 50 percent loss. The development in the OBX index during this period is illustrated in Figure 5.2.



Figure 5.2: The OBX index around the Russian crisis

The implied risk-neutral density functions are extracted at six different dates. Two of the distributions are estimated prior to the fall, at February 20, in the middle of the recovery after the Asian crisis, and at April 23, around the top level of the OBX index. The two following distributions are extracted during the fall, at July 24 and August 21. The next distribution is estimated at October 23, around the bottom level, while the final distribution is estimated six

months later, at April 23, after a modest recovery in the market. The implied distributions are presented in Figure 5.3, and the summary statistics are given in Table 5.1.

	RMSE			Standard deviation			Skewness		Probability-7.5% fall*		
	SPLINE	DLN	SLN	SPLINE	DLN	SLN	SPLINE	DLN	SPLINE	DLN	SLN
February 20, 1998	0.331	0.315	1.336	0.065	0.065	0.060	-1.085	-1.118	11.4	10.6	10.4
April 23, 1998	0.334	0.321	0.910	0.060	0.062	0.057	-0.690	-0.917	10.6	8.9	9.1
July 24, 1998	0.249	0.244	0.421	0.054	0.059	0.052	-0.411	-1.103	9.0	7.3	9.1
August 21, 1998	0.520	0.467	0.444	0.084	0.101	0.083	0.085	-0.810	18.4	16.7	18.9
October 23, 1998	0.429	0.429	1.170	0.107	0.108	0.102	-0.506	-0.534	23.5	23.1	23.8
April 23 1999	0.443	0 444	0.837	0.064	0.064	0.062	-0.532	-0.574	12.4	12.1	11.0

*Implied risk-neutral probability of a fall in the OBX index of more than 7.5 percent relative to the futures price the next four weeks. The other variable definitions are given in Section 4.4.

Table 5.1: Summary statistics for implied risk-neutral density functions extracted around the Russian crisis

By studying the DLN and SPLINE distributions extracted at February 20, we see that they have a pronounced left tail and differ significantly from the distribution obtained by the single lognormal model. This is confirmed by Table 5.1 which reports a skewness of about -1 for both the SPLINE and the DLN method. The large negative skewness may indicate that the market is not set at rest after the rapid decline caused by the Asian crisis, and is worried for a new decline in the stock market. Note that even if the DLN and SPLINE distributions exhibit large negative skewness, the implied probability of a fall in the OBX index is quite similar for all three methods.

By comparing the distributions at February 20 and April 23, only minor changes in the shape of the implied density functions are observed. Also when moving two months ahead, to July 24, the general picture remains the same. There seems to be no significant changes in the market expectations during this period. The standard deviation and the implied probability are only slightly reduced. The skewness reported by the DLN method stays almost constant, but the skewness reported by the SPLINE method is more than halved since February 20. However, if we compare the distributions visually, we see that the overall shape is about the same, and not too different from the SLN distribution. The large skewness reported by the DLN method is caused by a "thin" left tail. Anyway, the general impression from the implied density functions is that the uncertainty in the market seems to be slightly reduced. This indicates that the market participants are not expecting the dramatic decline following the next month.



Figure 5.3: Implied (logarithmic) return distributions extracted around the Russian crisis

At August 21, after a market fall of about 20 percent, the shape of the implied density functions are clearly changed. The standard deviations of the distributions are increased by roughly 60-70 percent, and the probability of a fall is more than doubled. This indicates an increased anxiety among market participants for the future stock market development. Notice that the shape of the SPLINE distribution and the DLN distribution this time is quite different. The SPLINE distribution is almost normal, while the DLN distribution is clearly negatively skewed. So according to the DLN distribution, the market participants are most worried of the downside risk, while according to the SPLINE method, the risk is more symmetrical.

At October 23, the DLN and SPLINE distributions are more harmonized. At this time, the OBX index has experienced its first "major" rise since the collapse started. We might therefore perhaps expect a more positive attitude to the future development. But both the standard deviation and the probability of a fall increase, and the skewness is still highly negative. This suggests that the market is still anxious and not convinced that the turbulence is over. However, if we move about six months ahead and look at the density functions extracted at April 24, the market participants seem to be less nervous. This is in line with what we might expect after the modest market recovery experienced the previous months. The standard deviation is reduced by about 40 percent, but it is still at the same level as observed in the beginning of the decline in the end of July. The implied probability of a fall in the OBX index is now in fact higher. Also the negative skewness is quite large. The high level of market uncertainty reflected by the implied risk-neutral density functions may appear a little bit strange when considering the future development in the OBX index. A possible explanation may be that the market participants in the Norwegian option market have become more worried for new "crashes" after experiencing the dramatic decline during the Russian crisis.

If we compare the implied risk-neutral density functions extracted using the smoothed implied volatility smile method and the double lognormal method with the distributions obtained by the single lognormal model, we see that there are obvious differences. Both the DLN and the SPLINE distributions (for return) are generally characterized by negative skewness, which the single lognormal model is not capable of representing. However, notice that the estimates of standard deviation and implied probability are quite similar for the three methods. By comparing also the root mean squared errors (RMSE) between observed option prices and implied option prices, we see that the poorest results are obtained for the single lognormal model. This indicates that OBX options are generally not priced according to the lognormal property of the Black-Scholes model. It also illustrates the weakness of using a single lognormal model to extract the implied risk-neutral density functions. A more comprehensive comparison of the various methods is presented in the next section.

It is worth noting that almost all of the implied risk-neutral density functions in this example exhibit negative skewness (using SPLINE and DLN). One reason, of course, may be that the market participants attach a high probability to a sharp decline in the stock market. Another possible explanation is that it may reflect the need for portfolio insurance⁶. Investors protect their portfolios against large downwards movement in the stock market by buying out-of-themoney put options. The increased demand drives up the prices of these option contracts, which is reflected as negative skewness in the implied risk-neutral density functions. Also recall that when deriving the implied risk-neutral density function we are assuming that market participants can hedge their positions perfectly. This assumption may not necessary hold. For example, in times of high market volatility, traders may be unwilling to write option contracts that provide insurance against large price falls. Market participants are therefore potentially more exposed to major declines in the stock market than to major rises, and may be willing to pay a greater premium for the downside insurance. Thus, a "neutral" implied distribution is not necessarily symmetric, but more likely negatively skewed.

5.3 Comparing Estimation Methods

In this section, the relative performance of the smoothed implied volatility smile method, the double lognormal model and the single lognormal is compared. First, the in-sample pricing errors are analysed by studying the differences between observed option prices and implied option prices from the various models. Then, a comparison of the summary statistics is performed. The analyses are based on time-series of implied risk-neutral density functions extracted each month from December 1995 until December 2001 with four weeks of remaining time to maturity, i.e. totally 73 distributions.

5.3.1 Pricing Errors

An important question when comparing different methods for extracting implied risk-neutral density functions is how good the different methods fit the observed option data. A comparison of the average root mean squared errors (RMSE) between observed option prices and implied option prices from the different models are presented in Table 5.2 (in NOK). The entries in Table 5.2 are obtained by first estimating the root mean squared error for each implied distribution, and then, the mean and standard deviation of the root mean squared errors for each year and for the whole period are calculated⁷.

The results for three slightly different variants of the SPLINE technique are presented. The differences are related to the extrapolation of the spline function outside the range of observations in implied volatility/delta space. SPLINE-1 uses the first and last polynomial to extend the spline

⁶See Grossman and Zhou [1996].

⁷Each year include the implied distributions extracted for all contracts expiring that year, except 2001, which also includes the January 2002 contract.

	SPLINE-1	SPLINE-2	SPLINE-3	DLN	SLN
1996	0.500(0.225)	0.500(0.224)	$0.617\ (0.173)$	$0.335\ (0.146)$	$0.480\ (0.173)$
1997	$0.422 \ (0.166)$	$0.422 \ (0.166)$	$0.569\ (0.088)$	$0.389\ (0.155)$	$0.488\ (0.214)$
1998	0.452(0.244)	$0.451 \ (0.244)$	$0.679\ (0.291)$	$0.429\ (0.219)$	0.849(0.348)
1999	$0.361\ (0.138)$	$0.360\ (0.137)$	$0.653\ (0.270)$	$0.346\ (0.144)$	0.869(0.233)
2000	$0.589\ (0.310)$	$0.589\ (0.309)$	1.034(0.448)	$0.576\ (0.313)$	$0.876\ (0.288)$
2001	0.468(0.202)	$0.467 \ (0.200)$	0.870(0.529)	$0.445\ (0.208)$	$0.664\ (0.190)$
1996 - 2002	$0.467 \ (0.225)$	$0.465\ (0.224)$	$0.739\ (0.366)$	$0.420\ (0.215)$	0.704(0.294)
The entries are t	he average root me	an squared errors ((RMSE) between o	bserved option pric	es and implied opt

are given in parentheses. RMSE is calculated as described in Section 4.4. The calculations are based on implied risk-neutral density functions extracted each month from December 1995 until December 2001 with four weeks of remaining time to maturity.

Table 5.2: Average root mean squared errors (RMSE) between observed option prices and implied option prices (standard deviation in parentheses)

function left and right, respectively, SPLINE-2 threats the spline function as linear outside the range of observation, while in SPLINE-3, lognormal distributions are fitted at the tails of the distribution. The different approaches are explained in detail in Section 4.3.3.

By comparing the different variants of the SPLINE technique, we see that the average pricing errors are about the same for SPLINE-1 and SPLINE-2, lying in the range 0.36-0.59. But surprisingly, the average errors for SPLINE-3 are considerably higher, varying from 0.57 to 1.03. This means that the tails of the implied distributions produced by SPLINE-1 and SPLINE-2 are almost similar, but quite different from the ones obtained by SPLINE-3. It also implies that SPLINE-1 and SPLINE-2 fit the option data better than SPLINE-3.



Figure 5.4: Implied risk-neutral density functions extracted by the different SPLINE methodologies at September 22, 2000

Figure 5.4 illustrates the difference in shape of the implied risk-neutral density functions that may occur between the different SPLINE methodologies. We see that the distributions obtained by SPLINE-1 and SPLINE-2 are about identical, while the distribution obtained using SPLINE-3 differs significantly from the other two. This occurs especially if the range of observable strike prices is small, as in this example. The small range of actively traded strike prices is a general problem in the Norwegian option market. For the time period studied, the probability mass corresponding to the observed range of strike prices varies between 40 to 90 percent, with a mean value of 65 percent. Thus, the non-observable part of the implied risk-neutral density functions is quite large. In a more liquid option market, I suppose the differences between SPLINE-3 and the other two approaches would have been significantly reduced. The small differences in average errors between SPLINE-1 and SPLINE-2 indicate that using the last polynomials at the ends to extrapolate the spline function or using a linear extrapolation is about equivalent. This is because the ends of the observable part of the volatility smiles are roughly linear, and that the extrapolation area in the implied volatility/delta space is relatively small.

If we compare the pricing errors for the SPLINE method, the double lognormal model and single lognormal model, we see that the double lognormal model fits the data slightly better than SPLINE-1 and SPLINE-2, which again fit the data considerably better than the single lognormal model, except for 1996. But rather surprisingly, for the whole period, SPLINE-3 performs worst. It only outperforms the single lognormal in 1998 and 1999. This means that the method of fitting lognormal distribution at the tails of the distribution, as implemented in this study, is not a suitable technique for modeling the tails of implied risk-neutral density functions extracted from OBX options.

The results in Table 5.2 suggest that market participants are not pricing options according to the Black-Scholes model, and that a single lognormal model is a too simple representation of the underlying price distribution implied by the market. The differences in average pricing errors between the double lognormal model and the SPLINE method (SPLINE-1 and SPLINE-2) may not be too surprising. In the double lognormal method, the distribution is obtained by minimized the pricing errors directly. In the smoothed implied volatility smile method, the implied volatilities, and not the prices, are approximated. This may led to a slightly more inaccurate representation of the option prices.

Another question of interest if how the mean of the implied risk-neutral density function deviates from the observed futures price. Theoretically, these are equivalent. When extracting the single lognormal distribution and the double lognormal distribution, the squared deviation between the future price and the mean of the distribution is added to the minimization problem ⁸. By not implementing this as a restriction, a meaningful comparison can be made between the implied mean and the observed futures price. For the smoothed implied volatility smile method, the futures price only enters when estimating implied volatilities and delta values. The average errors (absolute values) between the implied means and the corresponding futures prices for the

⁸See Section 4.2.

whole period are given in Table 5.3.

	SPLINE-1	SPLINE-2	SPLINE-3	DLN	SLN	
In NOK:	0.0532(0.1064)	$0.0531 \ (0.1123)$	2.0515(1.1863)	$0.1030\ (0.2616)$	0.5130(0.3966)	
In percent:	0.0021	0.0024	0.3222	0.0290	0.0817	
The entries are	the average errors (absolute values), in NO	K and in percent, b	etween observed future	s prices and implied means	. Standard

deviations are given in parentheses. The calculations are based on implied risk-neutral density functions extracted each month from December 1995 until December 2001 with four weeks of remaining time to maturity.

Table 5.3: Average absolute errors between implied means and observed futures prices (standard deviation in parentheses)

The results in Table 5.3 are promising. The average absolute error in percent between the implied means and the observed futures prices is roughly 0.002 for SPLINE-1 and SPLINE-2, 0.03 for the double lognormal model, and 0.08 for the single lognormal model. The results also confirm the relative weakness of SPLINE-3. An average absolute error around 0.3 percent is considerably higher than for the other models. Overall, the deviations between the implied mean and the futures price are negligible.

5.3.2 Summary Statistics

When comparing the various methods of extracting risk-neutral density functions, it is also of great interest to analyse the differences between the summary statistics of the implied distributions obtained by the different methods. I this section I perform a simple comparison of the basic summary statistics presented in Section 4.4. The comparison is based on a visual inspection of time-series plots of various summary statistics. I have also calculated correlation coefficients between the time-series obtained by the three different methods to better illustrate the relative strength of the relationships. The summary statistics included in the comparison are standard deviation, skewness, Pearson median skewness, excess kurtosis, and the two probability measures, the skewness-parameter and the uncertainty-parameter. Standard deviation, skewness, Pearson median skewness-parameter and the uncertainty-parameter are calculated from the implied *price* distribution. First, I will compare the various SPLINE methodologies. Then, the SPLINE method is compared with the double lognormal model and the single lognormal model.

Plots of the time-series and the estimated correlation coefficients are presented in Appendix A. By studying the plots in Figure 1 and Figure 2, the results in the previous subsection are strengthened. SPLINE-1 and SPLINE-2 produce about identical time-series for all the summary statistics. It is almost impossible to distinguish the time-series visually. The time-series obtained for SPLINE-3 deviate more or less from the other two. Not surprisingly, the deviations are larger for higher moments such as skewness and kurtosis. As pointed out earlier, these measures are very sensitive to the tails of the distributions. Large differences are also observed for the skewness-parameter. On the other hand, the time-series of standard deviation and the uncertainty-parameter are almost identical to the ones obtained for SPLINE-1 and SPLINE-2. The scatter plots in Figure 3 in Appendix A illustrate the differences between SPLINE-2 and SPLINE-3. In the scatter diagrams for standard deviation and the uncertainty-parameter, we see that the points are approximately lying on a 45 degree line through origo. This implies that the time-series of these measures are quite similar for the two methods. For the other summary statistics, the dispersion around the 45 degree line is significantly higher. Thus, the differences between these time-series are larger.

The correlation coefficients between the summary statistics obtained for the different SPLINE methodologies in Table 1 in Appendix A confirm the above observations. Between SPLINE-1 and SPLINE-2, all the correlation coefficients lie in the range 0.977-0.9995. The highest correlation is observed for standard deviation, and the lowest for the skewness-parameter. These large correlation coefficients imply that the implied risk-neutral density functions extracted using SPLINE-1 and SPLINE-2 are about identical. By studying the correlation coefficients for SPLINE-3, the results are more mixed. The correlation with SPLINE-1/SPLINE-2 is above 0.99 for the standard deviation and the uncertainty parameter, while it is down to roughly 0.80 for the skewness-parameter.

This simple comparison of the various SPLINE methodologies suggests that SPLINE-1 and SPLINE-2 produce about identical results, while some differences are observed between SPLINE-1/SPLINE-2 and SPLINE-3. The similarities in the time-series of standard deviation and the uncertainty-parameter imply that the dispersion of the distributions obtained by the various SPLINE methodologies are rather similar. On the other hand, if we consider the skewness measures, the differences between SPLINE-1/SPLINE-2 and SPLINE-3 increase. Thus, the deviations between SPLINE-1/SPLINE-2 and SPLINE-3 are mainly related to differences in the asymmetry of the implied risk-neutral density functions. The results in this section are consistent with the results obtained for the pricing errors, and suggest that the tails of the distributions obtained by using SPLINE-3 are quite different from the ones obtained by using SPLINE-1 or SPLINE-2.

In the time-series plots in Figure 4 and Figure 5 in Appendix A, the smoothed implied volatily smile method is compared to the double lognormal model and the single lognormal model⁹. Since skewness and excess kurtosis is zero for the single lognormal model¹⁰, only the time-series of standard deviation and the uncertainty-parameter are included for this model.

By studying the time-series of standard deviation, we see that the series are fairly similar. Table 1 in Appendix A reports a correlation coefficient of 0.95 between SPLINE-2 and DLN, and

⁹SPLINE represents the results obtained for SPLINE-2.

¹⁰The implied (logarithmic) return distribution is normal for the single lognormal model.

0.99 between SPLINE-2 and SLN. The correlation between DLN and SLN is 0.95. Thus, the different models are rather consistent when considering standard deviation. The same holds for the uncertainty-parameter. This means that the dispersion of the implied distributions obtained by these three methods are quite similar. On the other hand, a comparison of the time-series of skewness and excess kurtosis (SPLINE and DLN) is not very promising. Table 1 reports a correlation coefficient of 0.52 for excess kurtosis and 0.36 for skewness (using SPLINE-2). Using Pearson median skewness in stead of the standard skewness measure, the correlation coefficient is increased to 0.71. The correlation for the skewness-parameter is 0.56 (using SPLINE-2). The scatter diagrams in Figure 6 in Appendix A illustrate the differences between SPLINE-2 and DLN. We see that except for standard deviation and the uncertainty-parameter, the deviations from the 45 degrees line through origo are fairly large. These results imply that the tails of the implied risk-neutral density functions extracted using the smoothed implied volatility method and the double lognormal method are quite different. However, it is not surprising that these two methods do not produce identical tails. Remember that the region of actively traded OBX option is rather small. Consequently, as discussed in Melick and Thomas [1998], there is an infinite variety of probability masses outside the lowest and highest available strike price that can be consistent with the observed option prices.

Notice the large excess kurtosis often observed for the double lognormal model. This occurs when one of the two lognormal components has a large standard deviation, which is reflected as long thin tails in the final density. Such results are of course questionable. Since we have no actively traded options in these regions of the density function, the long thin tails are solely a result of the parametric nature of the double lognormal method.

The measures of skewness and kurtosis reported by the smoothed implied volatility smile method and the double lognormal method show that there are evident deviations from lognormality. In the period from December 1995 to December 2001, the implied risk-neutral density functions (for return) are characterized by negative skewness and positive excess kurtosis. This confirms the inadequacy of using a single lognormal model to extract the implied risk-neutral density functions. Models capable of representing skewness and excess kurtosis in returns should be applied when extracting risk-neutral density functions. Note that, as pointed out in Section 5.2.1, negative skewness does not necessarily mean that the market participants expect a large decline in the price of the underlying asset. It may also reflect the need for portfolio insurance among investors. Likewise, the high level of excess kurtosis implied by the risk-neutral density functions indicates not only that market participants in the Norwegian option market may expect large price changes, but also that they are willing to pay a higher premium for protection against such large price changes.

5.4 Studying the Effect of the Bid-Ask Spreads

The effect of the bid-ask spreads on the estimation of the implied risk-neutral density functions is analysed as outlined in Section 4.5. To reduce the extent of the analysis, the simulation is performed for only three implied density functions. These are extracted at August 27, September 24 and October 19, 2001, around the attack on World Trade Center September 11. Remaining time to maturity is four weeks. The development in the OBX index in this time period, and the dates at which the implied risk-neutral density functions are extracted, are illustrated in Figure 5.5. Only the smoothed implied volatility smile method¹¹ and the double lognormal method are considered. The purpose is not to perform a comprehensive analysis, but more to illustrate how uncertainty in option prices caused by the bid-ask spreads affects the estimation of the implied risk-neutral density functions. I will also show how this framework can be applied when assessing changes in market sentiment.



Figure 5.5: The OBX index around the attack on World Trade Center September 11, 2001

5.4.1 Stability of the Implied Distributions

The bid-ask quotations used as input in the simulations are given in Table 2 - Table 4 in Appendix B. The estimated confidence bands together with the corresponding implied (logarithmic) return distributions are presented in Appendix B, Figure 7 - Figure 12. The solid line represents the estimated density function, while the dotted lines are the upper and lower bounds of the confidence band. Recall that this band represents the confidence interval for the density function at each single point. Thus, the upper and lower bounds do not represent possible density functions. The lower bound necessarily integrates to less than unity, and the upper bound

¹¹Using a linear extrapolation of the spline function outside the observed range of deltas (SPLINE-2).

integrates to more than unity. The 90 percent confidence intervals for the summary statistic are presented in Table 5.4.

		\mathbf{DLN}			SPLINE*			
	5^{th} perct.	95^{th} perct.	Spread	5^{th} perct.	95^{th} perct.	Spread		
August 27:								
Mean	705.65	708.17	2.53	708.06	708.20	0.14		
Standard deviation	0.034	0.045	0.011	0.036	0.040	0.004		
Skewness	-0.71	0.36	1.07	-0.23	0.81	1.03		
Pearson median skewnes	-0.09	0.03	0.12	-0.05	0.11	0.16		
Excess kurtosis	-0.33	6.71	7.04	-0.26	1.49	1.75		
Skewness-parameter	-0.81	1.72	2.53	-0.71	4.65	5.36		
Uncertainty-parameter	2.10	8.96	6.86	3.23	7.82	4.59		
September 24:								
Mean	546.65	548.78	2.13	548.70	548.73	0.03		
Standard deviation	0.079	0.139	0.060	0.101	0.111	0.010		
Skewness	-2.27	0.00	2.27	-0.93	-0.33	0.60		
Pearson median skewnes	-0.21	-0.01	0.20	-0.22	-0.10	0.12		
Excess kurtosis	-0.01	7.78	7.79	0.17	0.95	0.77		
Skewness-parameter	-4.10	8.94	13.03	-2.70	1.85	4.56		
Uncertainty-parameter	32.14	44.40	12.27	38.89	45.03	6.14		
October 19:								
Mean	555.96	556.25	0.29	555.96	555.99	0.04		
Standard deviation	0.076	0.112	0.036	0.077	0.093	0.016		
Skewness	-2.21	-0.29	1.92	-1.18	-0.31	0.87		
Pearson median skewnes	-0.25	-0.11	0.13	-0.24	-0.09	0.15		
Excess kurtosis	0.09	7.35	7.26	0.02	1.56	1.54		
Skewness-parameter	-5.67	2.68	8.35	-3.10	-0.84	2.26		
Uncertainty-parameter	22.35	30.16	7.81	30.09	33.02	2.93		

The confidence intervals at each date are obtained by first simulating 500 price series by drawing uniformly from the bid-ask spreads. Then, the corresponding implied distributions are extracted from the simulated option price series, and the summary statistics are estimated for each distribution. Finally, the 5^{th} percentile and the 95^{th} percentile are calculated for each sample of the various summary statistics. Mean, skewness-parameter and uncertainty-parameter are calculated from the price distributions, the other summary statistics are calculated from the (logarithmic) return distributions.

*Using a linear extrapolation of the spline function outside the observed range of deltas (SPLINE-2).

Table 5.4: 90 percent confidence intervals for summary statistics obtained from bid-ask simulations

The plots of the confidence bands in Appendix B illustrate visually the uncertainty introduced by the bid-ask spreads. The wider the band is, the greater is the uncertainty surrounding the estimated density function. By studying the plots, we see that the confidence bands obtained for the double lognormal method are wider than the corresponding bounds obtained for the smoothed implied volatility smile method, especially in the tails of the implied distributions. Table 5.4 confirms the expression from the plots. The confidence intervals for the summary statistics are, almost without exceptions, greater for the double lognormal method than for the smoothed implied volatility smile method. If we look at the confidence intervals for the mean values, we observe that they are negligible for the SPLINE method. They are a bit higher for the double lognormal method, but small compared to the size of the mean values. The dispersion of standard deviation is relatively higher. For the SPLINE method, the intervals are fairly small. The largest interval is stretching from 0.077 to 0.093. The dispersion for the double lognormal method is considerably higher. For example, at September 24, the minimum value of the standard deviation is 0.079, while the maximum value is 0.139, nearly doubled. The corresponding interval for the SPLINE method is only 0.101-0.111.

Moving to the higher order moments, the results are not very promising. Extremely large dispersions in skewness and excess kurtosis are observed, so large that the estimates are almost useless. The same pattern is observed for the skewness-parameter, while the results for Pearson median skewness are slightly better. The variation in the estimates of the uncertainty-parameter is considerably smaller, especially at September 24 and October 19 when the uncertainty-parameter is relatively large. The dispersions are especially large for the double lognormal method. By comparing the two methods, we see that the confidence intervals for skewness are nearly the same at August 27, while at September 24 and October 19 the confidence interval is about 4 and 2 times as large for the double lognormal method as for the SPLINE method, respectively. For kurtosis, the confidence intervals are 4, 10 and 4.5 times larger for the double lognormal method at the respective dates. The differences are considerably smaller for the skewness-parameter and the uncertainty-parameter. At worst, at October 19, the dispersion of the skewness-parameter is about 3.5 times higher for the double lognormal method than for the SPLINE method, and about 2.5 times higher for the uncertainty parameter. The smallest difference between the two methods is observed for Pearson median skewness. The maximum difference is observed at October 19, when the confidence interval obtained by the double lognormal method is about 1.5 times as large as for the smoothed implied volatility smile method. Recall that this measure is less sensitive to the tails of the implied density function compared to the standard skewness measure.

Notice that even though the confidence bands plotted in Appendix B become quite narrow as we move towards the tails of the distributions, the relative size of the confidence bands increases. This can be illustrated by plotting the spread (confidence interval) across various n^{th} probability percentiles. First, the n^{th} probability percentile is calculated for each of the simulated density functions. From this sample, the median, the 95th percentile, and the 5th percentile is calculated. I then estimate the percentage deviation between the 95th percentile and the median, and between the 5th percentile and the median. This is performed for various probability percentiles. Figure 5.6 shows the dispersion across probability percentiles at September 24. The distance between the cross and the pluss sign can be thought of as the 90 percent confidence interval (in percent) for the respective probability percentile obtained by the SPLINE method. The distance between the upward triangle and the downward triangle is the corresponding confidence interval estimated by the double lognormal method.



Figure 5.6: Dispersion across probability percentiles at September 24, 2001

From Figure 5.6 we clearly see how the dispersion in the estimates of the probability percentiles increases as we move towards the tails of the distribution. As expected, the dispersion for the double lognormal method is generally greater than for the smoothed implied volatility smile method. The added uncertainty in the tails of the implied risk-neutral density functions explains the large confidence intervals reported in Table 5.4 for summary statistics sensitive to the tails of the implied distributions. Some of this uncertainty may be related to the limited range of strike prices observed for OBX options. As illustrated in Figure 5.7, the non-observable parts of the implied risk-neutral density functions are in this example quite large. Thus, as emphasized in Section 4.4, it is not surprising that there are differences in the tails of the distributions obtained by the smoothed implied volatility smile method and the double lognormal method, and that great disturbances in option prices may lead to large variations in the tails within the various estimation methods.

Why does the double lognormal method appear to be less stable than the smoothed implied volatility smile method? One obvious reason is that the double lognormal method is more sensitive to disturbances in option prices since they are directly used to estimate the density function. A local price change may affect the entire distribution. In the smoothed implied volatility smile method, the option prices are transformed to implied volatilities and a smoothed curve is fitted to the observations. The impact of disturbances in option prices is therefore less. Notice that by reducing the smoothing parameter, a greater part of the disturbances is smoothed away, which increases the stability of the smoothed implied volatility smile method. In the present calculations, I have used the same smoothing parameter for all the simulated distributions. The parameter is set equal to the one used for the unperturbed distribution, which is in disfavor of the smoothed implied volatility smile method. Due to the relatively large disturbances introduced in the option prices, using the original smoothing parameter may


Figure 5.7: Regions of observable strike prices at August 27, September 24 and October 19, 2001.

give less smooth distributions, and thus, reduced stability. The ideal solution is to choose an "optimal" smoothing parameter for each of the simulated distributions. But as pointed out in Section 4.3.2, a procedure for choosing the "optimal" smoothing parameter is not implemented, so the smoothing parameter had to be picked manually. In this example, with totally 3000 simulated distributions, this is impractical.

The results in this section show that the uncertainty introduced by the bid-ask spreads is quite large, especially when considering the tails of the implied risk-neutral density functions. Some of this uncertainty may also be due to the limited range of strike prices available in the Norwegian option market. This suggests that one should not place to much confidence in higher order summary statistics and in probability measures including the tails of the implied distributions. The analysis also illustrates that the implied risk-neutral density functions obtained by the double lognormal method are extremely sensitive to disturbances in option prices. The smoothed implied volatility smile method appears to be more robust to such disturbances. These findings are consistent with the more comprehensive studies by Bliss and Panigirtzoglou [1999] and Cooper [1999]¹². Both studies conclude that based on accuracy and stability of the estimated summary statistics, the smoothed implied volatility smile method outperforms the double lognormal method.

5.4.2 Assessing Changes in the Market Sentiment

So far, the discussion of the confidence bands has been rather technical. In this section, I will discuss possible economic, or practical, aspects of the confidence bands.

 $^{^{12}\}mathrm{Discussed}$ in Section 3.4.2.

As described in Section 3.4.2, Andersson and Lomakka [2001] derive confidence bands by perturbing the theoretical option prices by historical or actual price errors, and suggest to use them to decide whether changes in the implied risk-neutral density functions are statistical significant or just noise. This is important in a practical point of view, for example when assessing changes in market expectations due to specific economic events. Andersson and Lomakka [2001] classify an event as insignificant in a statistical sense if the density estimated after the event falls within the confidence band derived from the pre-event distribution, and significant if it falls outside. In the present analysis, I assume that the major uncertainty in the estimated distributions is related to using an average of the closing bid and ask quotation as a proxy for the option price. So if the post-event distribution falls outside the confidence band estimated for the pre-event distribution, it means that the option prices now are outside the previous bid-ask quotations. Thus, the changes in the option prices can be thought of as significant on a 90 percent level.

As an example of this approach, I will consider whether the market expectations implied from OBX option prices changed significantly around the attack on World Trade Center. The confidence bands (dotted lines) and the implied risk-neutral density functions (solid lines) extracted the following month using the double lognormal method are illustrated in Figure 5.8. The similar results for the smoothed implied volatility smile method are presented in Appendix B, Figure 13, Figure 15 and Figure 17.



Figure 5.8: Confidence bands (dotted lines) and implied distributions (solid lines) extracted the following month

From the first plot in Figure 5.8, we see that the implied risk-neutral density function extracted from OBX options is clearly changed from August 27 to September 24. The latter distribution is much wider and reflects an increased uncertainty for the future development in the Norwegian stock market. This suggests that the market expectations are significantly revised after the attack on World Trade Center. By comparing the confidence bands extracted at September 24, and the implied risk-neutral density function at October 19, we observe that the changes in the market expectations during this period are considerably smaller, but still significant. The uncertainty among the market participants seems to be reduced. From October 19 to November 23, the changes are almost negligible. Only very small portions of the implied distribution falls outside the confidence band extracted the previous month. This suggests that the market expectations implied from OBX options are only slightly changed in the period from October 19 to November 23. But also in this period, the changes are significant on a 90 percent level.

As pointed out in Section 3.2, it is important to distinguish the risk-neutral distribution from the market's subjective probability distribution. The two distributions are only equivalent if investors act as if they are risk-neutral. But as Rubinstein [1994] shows, if we assume that the representative investor has constant relative risk aversion, the "true" distribution will shift to the right, while the shape is about unchanged. Thus, assessing changes in the implied risk-neutral density functions may give us valuable information about alterations in market expectations. However, this requires that the degree of risk aversion remains constant in the time period studied. Hence, comparing implied distributions on a monthly basis as illustrated above, may be questionable since the general level of risk aversion among investors may have changed. A more ideal analysis would be to consider changes during a shorter time horizon, for example a day or even hours.

5.5 Are Implied Risk-Neutral Density Functions Useful as Leading Indicators?

Option prices are assumed to contain unique information about how market participants assess the likelihood of different outcomes for future market prices. In this section, I will analyse the information contained in OBX option prices, and study whether properties of implied riskneutral density functions are useful as leading indicators in the Norwegian stock market. I will start the analysis by studying the relationships between the OBX index and measures that we intuitively might expect contain information useful for predicting future market prices. Then, more formal tests are considered.

5.5.1 Studying Possible Leading Indicators

Figure 5.9 shows the time-series of the implied risk-neutral probability of a fall in the OBX index of 7.5 percent or more relative to the futures price the next four weeks, together with the development in the OBX index. The implied probabilities are presented for both the double

lognormal method (DLN) and the smoothed implied volatility smile method (SPLINE)¹³, and are plotted at the observation dates. As previously, the implied distributions are extracted each month from December 1995 to December 2001 with four weeks of remaining time to maturity.



Figure 5.9: Implied risk-neutral probability of a fall in the OBX index of 7.5 percent or more relative to the futures price the next four weeks vs OBX index (equal to Figure 20 in Appendix C)

The general impression from Figure 5.9 is that market expectations are revised after a major drop occurs, and not prior to the fall. The peak values for the implied probabilities are generally associated with a recent decline in the stock market. For example, the implied probability rose sharply as the market fell in connection with the Russian crisis¹⁴. There was only a small upward trend in the implied probability prior to this crisis, which can not be interpreted as the market expected the dramatic decline experienced during the Russian crisis. The same jump in the implied probability is observed after the attack on World Trade Center.

Following the Russian crisis, the Norwegian stock market experienced a long and steady upward movement. From October 1998 to the beginning of September 2000, the OBX index increased from about 400 to 900, a rise of roughly 125 percent. During this period, there is a negative trend in the implied probability. This is perhaps contrary to what we might expect. As the market continued to grow, we would expect an increased anxiety for a decline in the OBX index. Internationally, there was increased focus in this period on the possibility that stock markets world wide, and in US in particular, were overvalued. In this situation we would expect increased demand for insurance (hedge) against this risk, which would result in a more

¹³Using a linear extrapolation of the spline function outside the observed range of deltas (SPLINE-2).

 $^{^{14}}$ See Figure 5.1 in Section 5.2 for an overview of the various crisis.

negatively skewed distribution and a rise in the implied probability of a fall. An interesting question is to what extent the same counterintuitive development took place in other option markets. The US market would be of special interest in this respect. One factor that might justify the development of the implied probability in the Norwegian market in this period was the rising oil price, which naturally would be considered as improving the prospects for the Norwegian economy and the stock market.

In Figure 19 in Appendix C, the time-series of implied standard deviation is plotted against the time-series of the OBX index. If we compare the plots of the implied probability and the implied standard deviation, we see that they are closely linked. The same conclusion can be drawn if we study the time-series of the uncertainty-parameter in Figure 21 in Appendix C. The impression from the plots is confirmed by the correlation coefficients presented in Table 5 and Table 6 in Appendix C. The correlation coefficient between implied standard deviation and implied probability is 0.98 for the smoothed implied volatility smile method, and 0.84 for the double lognormal method. Between the uncertainty-parameter and the implied probability the correlation is 0.98 (SPLINE) and 0.97 (DLN). Thus, the general pattern of the three time-series is quite similar. This implies that the general level of uncertainty in the market, measured as the spread of the implied distribution, is increasing substantially after a major drop in the stock market, and not prior to. It also means that the implied probability does not give a good representation of asymmetries within the implied risk-neutral density function. When the implied probability of a fall increases, the probability of a rise may also increase.

A measure that perhaps better captures asymmetries in market expectations is the skewnessparameter¹⁵. This measure represents the excess probability of a large rise relative to a large fall in the OBX index. Thus, when it is positive, the implied risk-neutral probability of a certain increase in the OBX index is greater than the probability of a similar decrease. When it is negative, the opposite is true. The time-series of the skewness-parameter and the OBX index is presented in Figure 5.10.

By studying Figure 5.10, there seem to be no intuitive relationships between the skewnessparameter and the OBX index. During the upward movements experienced prior to the Asian crisis and prior to the turbulence caused by the "collapse" in parts of the IT-industry in fall 2000, there are no specific trends in the skewness-parameter. This suggests that there were no signs of increased anxiety among market participants for a rapid decline in the stock market in these periods. Also, a large negative skewness-parameter is generally not associated with a subsequent fall in the stock market. However, note that large positive peak values of the skewness-parameter are often associated with a recent decline in the stock market. This may suggest that market participants often expect the stock market to recover relatively soon after a sudden decline.

 $^{^{15}\}mathrm{Must}$ be distinguished from skewness. Definition is given in Section 4.4.



Figure 5.10: Skewness-parameter vs OBX index (equal to Figure 22 in Appendix C)

An interesting observation from Figure 5.10 is that after the Russian crisis there is a significant fall in the general level of the skewness-parameter. The market participants now seem to be more worried about a major decline in the stock market than prior to the Russian crisis. As discussed earlier, this is not necessarily because they were expecting the market to fall. More likely is the increased negative skewness-parameter in this period due to increased risk aversion among investors.

The above discussion suggests that the information contained in implied risk-neutral density functions extracted from OBX options are not useful in predicting major declines in the Norwegian stock market. The expectations are generally revised after, and not prior to, a rapid decline in the market. As the market drops, the level of uncertainty increases. But often, also the implied risk-neutral probability of a large rise relative to a large fall increases. More formal tests of the information content of OBX option prices are considered next.

5.5.2 Realized versus Implied Volatility

The standard deviation calculated from the implied risk-neutral density function represents the market's opinion about the future volatility of the underlying asset. It is therefore of great interest to study whether the implied volatility can be used to predict the realized volatility.

Implied volatility and realized volatility of daily returns for the remaining time to maturity are calculated as described in Section 4.6.1 by equation (4.30) and (4.31). In Figure 5.11, monthly observations of predicted volatility are plotted against the corresponding realized volatility the following four weeks. As noticed in Section 5.3.2, there are only minor differences between the time-series of standard deviation calculated from the various models. Thus, only the results for the double lognormal model are presented.



Figure 5.11: Predicted volatility (DLN) and realized volatility

The impression from Figure 5.11 is that during calm periods in the stock market, the predicted volatility is a fairly good estimate of the future volatility. This is observed in 1996 through first half of 1997, and in the second half of 1999. In turbulent periods, on the other hand, quite large deviations are observed between predicted and realized volatility. In these periods, the predicted volatility is rather backward looking than forward looking. This is especially observed during the second half of 1998 through first half of 1999, and in the second half of 2001.

As presented in Section 4.6.1, the predictive power of implied volatility can be tested by the regression model in (4.32). But for the sample data used in the present analysis, the model appears not to be very robust. The error terms were autocorrelated and heteroscedastic, and not normally distributed. Inference based on the regression model is therefore not valid. Instead, to illustrate the relationship between the implied volatility and the future realized volatility, I have made a scatter plot of the two variables. For comparison, I have also made a scatter plot of realized volatility versus historical volatility the previous 30 days. The result, using the double lognormal model, is presented in Figure 5.12. The scatter plots for the smoothed implied volatility smile method and the single lognormal model are presented in Figure 23 and Figure 25 in Appendix C. Only minor differences are observed between the scatter plots.

Figure 5.12 shows that on average, there is almost a one to one relationship between predicted and realized volatility. But there are clearly large deviations from the regression line. With a R^2 measure of 0.393, it means that only roughly 40 percent of the total variation in the realized volatilities are explained by the implied volatilities. However, we see that the implied volatilities contain more information about future volatility than the historical volatilities do. The R^2



Figure 5.12: Scatter plots of realized volatility versus predicted volatility (using DLN) and realized volatility versus historical volatility (left panel equal to Figure 24 in Appendix C)

measure is significantly reduced, and a slope coefficient of 0.6 means that using the historical volatility is on average a more biased estimate of the future volatility. These results suggest that the volatility extracted from the implied risk-neutral density function provide some useful information in forecasting the future volatility of the OBX index.

5.5.3 Information Content of Asymmetries within the Implied Risk-Neutral Density Functions

As outlined in Section 4.6.2, a simple non-parametric sign test is proposed as a test of whether asymmetries within the implied risk-neutral density function can provide valuable information about future market prices. The hypothesis to be tested is whether these asymmetries can be used to explain the future deviation between the spot price at maturity (S_T) and the futures price at the observation date (F_0) . The proposed measures (Q) for predicting the deviation between the futures price and the future spot price are based on skewness, Pearson median skewness and the skewness-parameter as described in Section 4.6.2. These measures capture the asymmetry of the implied risk-neutral density function, and may contain valuable information about the market participants' assessment of the direction in the future development.

As preciously, the implied risk-neutral density functions are extracted each month from December 1995 to December 2001 with four weeks of remaining time to maturity. The deviation between the futures price at the observation date and the spot price at maturity is then calculated for each monthly observation. Q is estimated from the implied risk-neutral density functions as described in Section 4.6.2. This gives us a test sample of 73 observations for Q and $(S_T - F_0)$. Since the single lognormal model does not provide any information about asymmetries in market expectations, the tests are performed only for the double lognormal method and the smoothed implied volatility smile method.

The results from the non-parametric sign test are presented in Table 5.5. "Normal" level for all the measures are defined relative to the distribution of $(S_T - F_0)$ as described in Section 4.6.2 such that the number of positive and negative values for Q and $(S_T - F_0)$ are equally distributed¹⁶. Similar results using the mean and the median value as the "normal" level are given in Table 7 and Table 8 in Appendix C. The entries presented are the number of times the predictions of the model is correct. The numbers in parentheses are the significance levels from one-sided binomial tests with equal probability of success and failure. The null hypothesis of no predictive power implies that the number of correct predictions as a fraction of total observations should be less or equal to 0.5. I have also considered the ability of the models to predict upward movements $(S_T - F_0 > 0)$ and downwards movements $(S_T - F_0 < 0)$ separately.

	Skewness	Pearson	Skewness-
		median	parameter
		skewness	
SPLINE*:			
Upward and downward movement $(n=73)$	37~(0.408)	39(0.241)	41 (0.121)
Upward movement $(n=43)$	25~(0.111)	26~(0.063)	27 (0.033)
Downward movement $(n=30)$	12 (0.819)	$13 \ (0.708)$	14 (0.572)
DLN:			
Upward and downward movement (n=73)	33~(0.759)	35~(0.592)	41 (0.121)
Upward movement $(n=43)$	23 (0.271)	24 (0.180)	27 (0.033)
Downward movement $(n=30)$	$10 \ (0.951)$	$11 \ (0.890)$	14 (0.572)

The entries are the number of correct predictions in the sign test presented in Section 4.6.2. "Normal" level for all the measures (Q) are defined relative to the distribution of $(S_T - F_0)$ such that the number of positive and negative values for Q and $(S_T - F_0)$ are equally distributed. The numbers in parentheses are the significance levels from one-sided binomial tests with equal probability of success and failure. *Using a linear extrapolation of the spline function outside the observed range of deltas (SPLINE-2).

Table 5.5: Number of correct predictions in sign test (significance values in parentheses)

The results from Table 5.5 are not encouraging. There seems to be no predictive power in the the various skewness measures. Considering the ability of the models to predict both upward and downward movements, the hypothesis of no predictive power can not be rejected at a 5 percent significance level for any of the measures. Considering only upwards movement, the results are slightly better. For both methods, the value for the skewness-parameter is now significant at a 5 percent significance level. But a prediction record of 27 out of 43 is not very impressive. The results considering only downward movements are consequently poorer. There are only minor differences between the estimation methods and between the various skewness measures. The prediction results obtained by the smoothed implied volatility smile method are generally slightly better than the results obtained by the double lognormal method, and the

 $^{^{16}}S_T - F_0$ is greater than zero 43 times, and smaller 30 times.

skewness-parameter seems to contain some more information than skewness and Pearson median skewness. By studying the results in Table 7 and Table 8 in Appendix C, we see that the ability to predict future market prices are generally worse when using the mean or the median as the "normal" level.

5.5.4 Discussion of Results

The results in this section suggest that the information contained in implied risk-neutral density functions extracted from equity options on the OBX index is of little use in predicting future market prices. By first studying time-series plots of possible leading indicators extracted from the implied distributions, we saw that market expectations generally are revised after, and not prior to, a rapid decline in the stock market. As the market drops, the level of uncertainty increases. Often, the implied risk-neutral probability of a large rise relative to a large fall also increases. More formal tests confirm the lack of predictive power contained in OBX option prices. However, the standard deviation estimated from the implied distributions may provide some useful information in forecasting the future volatility of the OBX index. This is in line with the results in Navatte and Villa [2000] and Weinberg [2001]¹⁷. A positive finding is that the implied volatility seems to contain more information about future volatility than the historical volatility does. The various skewness measures, on the other hand, seem to contain no information about the direction of future outcomes. The lack of forecasting power in skewness is also found by Weinberg [2001] when testing the ability of implied skewness to predict realized sample skewness. Similar analysis by Navatte and Villa [2000] reports that implied skewness contain a substantial amount of information for future skewness. So the conclusions on the predictive power of skewness are mixed.

That OBX options prices contain little information about future market prices may not be a big surprise. There are many obvious reasons for that. One reason may simply be that the views represented by market participants are generally not consistent with the future development. If we are more optimistic about the forward looking nature of market agents, and assume that market expectations contain some valuable information about future market prices, the low liquidity in the Norwegian option market may be a reason for the poor results obtained in this study. The number of traded option contracts and the range of available strike prices are relatively small, which limit the information content of option prices. One may of course argue that if those who operate in the Norwegian option market are few, but sophisticated, option prices will still contain a rich source of information. However, as we have seen in this study, the uncertainty related to the tails of the implied risk-neutral density functions is quite large, and using information from this region of the distributions is questionable. Another possible reason for the lack of predictive power is that the density functions extracted are risk-neutral. Market

 $^{^{17}}$ See Section 3.5.

participants may be willing to pay a premium to insure against price movements. The markets assessment of the true statistical probability of large price movements are therefore masked by the need for portfolio insurance. A final possibility, and probably the most important reason, is that the market is efficient. The expectations and information of all market participants are constantly fully reflected in stock market prices. Using information in option prices that is already incorporated in stock market prices to predict the future stock market development is therefore not possible.

Chapter 6

Summary and Conclusions

Option prices are assumed to contain unique information about how market participants assess the likelihood of different outcomes for future market prices. A popular way of gaining this information is by estimating implied risk-neutral density functions. One objective of this study is to implement techniques for estimating implied risk-neutral density functions, and to study the uncertainty surrounding the estimation of implied density functions in the Norwegian option market. The second objective is to analyse the potential value of the information embedded in OBX option prices. The aim is to gain a better understanding of whether properties of implied risk-neutral density functions can be used as leading indicators in the Norwegian stock market.

Three different methods for estimating implied risk-neutral density functions are implemented. The simplest model assumes that the underlying price distribution used by market participants when pricing options is a lognormal distribution (the single lognormal model). This model is equivalent to the famous Black-Scholes model. The second method, the double lognormal method, assumes that the underlying distribution is a combination of two lognormal distributions. For both methods, the parameters of the implied distribution are estimated by minimizing the squared deviation between observed option prices and theoretical model prices. The last approach is quite different. By interpolating the implied volatility across delta values using a smoothing spline function, transforming it to the implied volatility/strike price space, and substituting it in the Black-Scholes model, the call price can be expressed as a function of strike price. Using a famous result known as the Breeden-Litzenberger formula, we can easily obtain an expression for the risk-neutral density function. The method is called the smoothed implied volatility smile method. Due to the non-parametric nature of this approach, there is no obvious way of modeling the tails of the distributions. Thus, three alternative ways of modeling the tail are implemented. The first alternative extrapolates the spline function outside the range of observations in implied volatility/delta space by using the first and last polynomial to extend it left and right, respectively. The second treats the spline function as linear outside the range of observations, while in the last alternative, lognormal distributions are fitted at the tails of the distribution.

The analysis starts with a comparison of the different estimation methods. The comparison is based on monthly estimations of implied risk-neutral density functions in the period from December 1995 to December 2001. Remaining time to maturity for the options is four weeks. Studying the in-sample pricing error shows that the double lognormal method fits the option price data slightly better than the smoothed implied volatility smile method. Both methods fit the data considerably better than the single lognormal method. Comparing the three slightly different variations of the smoothed implied volatility smile method shows that using the end polynomials to extrapolate the spline function, or assuming that it is linear outside the range of observations, are about equivalent both when considering pricing errors and comparing the summary statistics of the implied distributions. The third approach, fitting lognormal distributions to the tails of the density function, appears not to be a very good solution the way it is implemented in this study. The pricing errors are relatively larger, and the summary statistics reveal that the tails obtained using this approach are quite different from the ones obtained using the other two approaches. The differences are probably a result of the small range of observable strike prices often observed for OBX options.

Comparing time-series of measures representing the spread of the implied distributions shows a high level of similarity among the estimation methods. On the other hand, studying various skewness measures and kurtosis, the differences increase. It confirms the inadequacy of using a single lognormal model to extract implied risk-neutral density functions. The implied return distributions obtained by the double lognormal method and the smoothed implied volatility smile method are characterized by negative skewness and positive excess kurtosis, features that the single lognormal model is not capable of representing. Comparing the double lognormal method and the smoothed implied volatility method shows that there are quite large differences between the tails of the implied distributions obtained by the two methods. A great part of these differences is probably due to the small region of actively traded OBX options. As the range of available strike prices gets smaller, the variety of probability messes outside the lowest and highest available strike price that can be consistent with the observed option prices increase.

The next part of the analysis considers how uncertainty in option prices affects the estimation of implied risk-neutral density functions. By using an average of the closing bid and ask quotation as a proxy for the option price, large bid-ask spreads may introduce a high level of uncertainty in the estimated distributions. To assess this uncertainty, new distributions are simulated by drawing new sets of option prices from the bid-ask spreads, and dispersions in summary statistics and in the shape of the distribution are studied. Only the double lognormal method and the smoothed implied volatility smile method are considered. The results shows that the uncertainty introduced by the bid-ask spreads are quite large, especially when considering higher order moments and tail probabilities. Sometimes the dispersion is so large that the estimates are almost useless. This suggests that one should not place too much confidence in higher order summary statistics and in probability measures including the tails of the implied distributions. The analysis also illustrates that the implied risk-neutral density functions obtained by the double lognormal method are extremely sensitive to disturbances in option prices. The smoothed implied volatility smile method appears to be more robust to such disturbances. As an extension of this analysis, I also show how this framework can be utilized when assessing changes in market expectations implied by the risk-neutral density functions.

The final section of the analysis studies whether properties of implied risk-neutral density functions extracted from OBX options are useful as leading indicators in the Norwegian stock market. Time-series of possible indicators derived from the implied distributions are studied first. The analysis suggests that in the period from December 1995 to December 2001, there are no indications that participants in the Norwegian option market were expecting the major declines occurring in this period. The general impression is that market expectations are revised after, and not prior to, a rapid decline in the stock market. More formal tests confirm the lack of predictive power contained in OBX option prices. However, the standard deviation estimated from the implied distributions is found to contain some useful information in forecasting the future volatility of the OBX index. The implied volatility seems to contain more information about future volatility than the historical volatility does. On the other hand, the ability of various skewness measures to predict the direction of future outcomes, seems to be poor. This is tested by a simple non-parametric sign test which considers the ability of the skewness measures to predict the future deviation between the spot price at maturity and the futures price at the observation date.

To sum up, this study shows that a single lognormal model is a too simple representation of the underlying price distribution employed by the market when pricing OBX options. Methods such as the double lognormal method and the smoothed implied volatility smile method, capable of representing skewness and excess kurtosis in returns, should be applied. The double lognormal methods appears to fit the option data slightly better than the smoothed implied volatility smile method, but is extremely sensitive to disturbances in option prices. The smoothed implied volatility method is more robust to such disturbances. The study also reveals that there is a high level of uncertainty surrounding the implied risk-neutral density functions extracted from OBX options. Uncertainty introduced by using an average of the closing bid and ask quotation as a proxy for the option price, and the small range of actively traded strike prices, suggest that we should not place too much confidence in estimates of higher order summary statistics and tail probabilities. The small range of actively traded strike prices is probably also a major reason for the differences often observed between the implied distributions obtained by the double lognormal method and the smoothed implied volatility method. Using information contained in OBX option prices in forecasting future market prices seems to be worthless. Some information about future volatility may be obtained, but not about the direction of future outcomes.

The high level of uncertainty surrounding the estimated risk-neutral density functions extracted from OBX options, and the fact that they seem to contain no valuable information in forecasting future market prices, may of course question the value of extracting implied distributions from OBX options. Certainly, it suggests that they should not be used for risk management purposes. For example, value-at-risk calculations based on tail probabilities may be very misleading. However, even if the implied distributions not are forward looking, they may still be useful for studying changes in the market sentiment. For example, for policy-making purposes, it may be of great interest to study sudden shifts in the implied distributions due to a political announcement or economic news. Such studies may give us valuable information about how the market reacts to specific announcements and news. This may be a possible direction for future studies on implied risk-neutral density functions in the Norwegian option market.

Bibliography

- Y. Aït-Sahalia and A. Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, 53(2):499–547, April 1998.
- M. Andersson and M. Lomakka. Evaluating implied RNDs by some new confidence interval estimation techniques. Working paper, Stockholm School of Economics, 2001.
- S. Aparicio and S. Hodges. Implied risk-neutral distribution: A comparison of estimation methods. Working paper, Warwick University, 1998.
- B. Bahra. Implied risk-neutral probability density functions from options prices: Theory and application. Working paper, Bank of England, 1997.
- D. S. Bates. The crash of '87: Was it expected? The evidence from option markets. *Journal* of Finance, 45(3):1009–1044, 1991.
- F. Black. The pricing of commodity contracts. Journal of Financial Economics, (3):167–179, March 1976.
- F. Black and M. Scholes. The pricing of options and corporate liabilities. Journal of Political Economy, (81):637–659, May-June 1973.
- R. R. Bliss and N. Panigirtzoglou. Testing the stability of implied probability density functions. Working paper, Bank of England, 1999.
- D. Breeden and R. Litzenberger. Prices of state-contingent claims implicit in option prices. Journal of Business, 51(4):621–651, 1978.
- P. W. Buchen and M. Kelly. The maximum entropy distribution of an asset inferred from option prices. Journal of Financial and Quantitative Analysis, 31(1):143–159, 1996.
- J. M. Campa, P. H. K. Chang, and R. L. Reider. Implied exchange rate distributions: evidence from OTC option markets. *Journal of International Money and Finance*, (17):117–160, 1998.

- J. Y. Campbell, A. W. Lo, and A. C. MacKinley. The Econometrics of Financial Markets. Princeton University Press, Princeton, New Jersey, 1997.
- N. Cooper. Testing techniques for estimating implied risk-neutral densities from the prices of european and american-style options. Mimeo, Bank of England, 1999.
- S. Coutant, E. Jondeau, and M. Rockinger. Reading PIBOR futures option smiles: The 1997 snap election. Working paper 54, Banque de France, 2000.
- P. Craven and G. Wahba. Smoothing noisy data with spline functions. Numerische Matematik, 31:377–403, 1979.
- G. Gemmill and A. Saflekos. How useful are implied distributions? Evidence from stock-index options. The Journal of Derivatives, 7(3):83–98, 2000.
- S. Grossman and Z. Zhou. Equilibrium analysis of portfolio insurance. Journal of Finance, 51 (4):1379–1403, 1996.
- S. L. Heston. A closed form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–344, 1993.
- J. C. Hull. *Options, Futures & Other Derivatives*. Prentice-Hall International, Englewood Cliffs, New Jersey, 4th edition edition, 2000.
- J. C. Jackwerth. Option-implied risk-neutral distributions and implied binomial trees: A literature review. *Journal of Derivatives*, 7(4):66–82, 1999.
- J. C. Jackwerth and M. Rubinstein. Recovering probability distributions from option prices. Journal of Finance, 51(5):1611–1631, 1996.
- R. Jarrow and A. Rudd. Approximate option valuation for arbitrary stochastic processes. Journal of Financial Economics, (10):347–369, 1982.
- E. Kreyszig. Advanced Engineering Mathematics. John Wiley & sons, Inc., 7th edition, 1993.
- M. Lomakka. Aktieindexoptioner som framtblickande indicator. *Penning- och Valutapolitik*, (3):68–78, 2001.
- D. J. MacManus. The information content of interest rate futures options. Working paper 99-15, Bank of Canada, 1999.
- D. B. Madan and F. Milne. Contingent claims valued and hedged by pricing and investing in a basis. *Mathematical Finance*, (4):223–245, 1994.

- A. M. Malz. Using option prices to estimate realignment probabilities in the European monetary system. *Journal of International Money and Finance*, 15(5):717–748, 1996.
- A. M. Malz. Estimating the probability distribution of the future exchange rate from options prices. *Journal of Derivatives*, 5(2):18–36, 1997.
- W. R. Melick and C. P. Thomas. Recovering an asset's implied pdf from options prices: an application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis*, 32(1):91–115, 1997.
- W. R. Melick and C. P. Thomas. Confidence intervals and constant maturity series for probability measures extracted from option prices. Paper presented at the Bank of Canada conference on Information Contained in Prices of Financial Assets, 1998.
- R. Merton. Theory of rational option pricing. The Bell Journal of Economics and Management Science, 4(1):141–183, 1973.
- P. Navatte and C. Villa. The information content of implied volatility, skewness and kurtosis: empirical evidence from long-term CAC 40 options. *European Financial Management*, 6 (1):41–56, 2000.
- S. Ross. Options and efficiency. Quarterly Journal of Economics, (90):75–89, 1976.
- M. Rubinstein. Alternative tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from august23, 1976 through august 31, 1978. Journal of Finance, 40(2):455–480, 1985.
- M. Rubinstein. Implied binomial trees. Journal of Finance, 49(3):771-818, 1994.
- B. Sherrick, S. Irwin, and D. Forster. Option-based evidence of the nonstationarity of expected S&P 500 futures price distributions. *Journal of Future Markets*, 12(3):275–290, 1992.
- B. Sherrick, S. Irwin, and D. Forster. An examination of option-implied S&P 500 futures price distributions. *Financial Review*, 31(3):667–694, 1996a.
- B. J. Sherrick, P. Garcia, and V. Tirapattur. Recovering probabilistic information from option markets: Test of distributional assumptions. *Journal of Futures Markets*, 16(5):545–560, 1996b.
- D. C. Shimko. Bounds of probability. *Risk*, 6(4):33–37, 1993.
- P. Söderlind. Market expectations in the UK before and after the ERM crisis. *Economica*, (67):1–18, 2000.

- P. Söderlind and L. Svensson. New techniques to extract market expectations from financial instruments. *Journal of Monetary Economics*, (40):384–429, 1997.
- S. A. Weinberg. Interpreting the volatility smile: an examination of the information content of option prices. International Finance Discussion Paper 706, Federal Reserve Board, 2001.





Standard deviation

Figure 1: Time-series of standard deviation, skewness and Pearson median skewness for the different SPLINE methodologies



Figure 2: Time-series of kurtosis, skewness-parameter and uncertainty-parameter for the different SPLINE methodologies



Figure 3: Scatter plot of SPLINE-2 and SPLINE-3



Figure 4: Time-series of standard deviation, skewness and Pearson median skewness



Figure 5: Time-series of excess kurtosis, skewness-parameter and uncertainty-parameter



Figure 6: Scatter plot of SPLINE-2 and DLN

	STAP	DARD D	EVIATION	ſ			SK	EWNESS		
	SPLINE	SPLINE	SPLINE	DLN	SLN		SPLINE	SPLINE	SPLINE	DLN
	-1	-2	-3				-1	-2	-3	
SPLINE-1	1	0.9995	0.9957	0.9466	0.9855	SPLINE-1	1	0.9917	0.9279	0.3077
SPLINE-2	0.9995	1	0.9956	0.9527	0.9863	SPLINE-2	0.9917	1	0.9258	0.3570
SPLINE-3	0.9957	0.9956	1	0.9329	0.9813	SPLINE-3	0.9279	0.9258	1	0.2335
DLN	0.9466	0.9527	0.9329	1	0.9515	DLN	0.3077	0.3570	0.2335	1
SLN	0.9855	0.9863	0.9813	0.9515	1					
	PEARSO	N MEDIA	N SKEWN	IESS			EXCES	SS KURTC	SIS	
	SPLINE	SPLINE	SPLINE	DLN			SPLINE	SPLINE	SPLINE	DLN
	-1	-2	-3				-1	-2	-3	
SPLINE-1	1	0.9874	0.9451	0.6393	-	SPLINE-1	1	0.9985	0.9707	0.4978
SPLINE-2	0.9874	1	0.9340	0.7054		SPLINE-2	0.9985	1	0.9696	0.5183
SPLINE-3	0.9451	0.9340	1	0.6653		SPLINE-3	0.9707	0.9696	1	0.4292
DLN	0.6393	0.7054	0.6653	1		DLN	0.4978	0.5183	0.4292	1
	UNCER	TAINTY-F	PARAMET	ER			SKEWNE	SS-PARAN	IETER	
	SPLINE	SPLINE	SPLINE	DLN	SLN		SPLINE	SPLINE	SPLINE	DLN
	-1	-2	-3				-1	-2	-3	
SPLINE-1	1.0000	0.9990	0.9955	0.9598	0.9797	SPLINE-1	1.0000	0.9771	0.7989	0.4933
SPLINE-2	0.9990	1.0000	0.9966	0.9642	0.9806	SPLINE-2	0.9771	1.0000	0.7810	0.5576
SPLINE-3	0.9955	0.9966	1.0000	0.9693	0.9858	SPLINE-3	0.7989	0.7810	1.0000	0.3936
DLN	0.9598	0.9642	0.9693	1.0000	0.9690	DLN	0.4933	0.5576	0.3936	1.0000
SLN	0.9797	0.9806	0.9858	0.9690	1.0000					

 Table 1: Correlation coefficients for standard deviation, skewness, Pearson median skewness,

 excess kurtosis, skewness-parameter and uncertainty-parameter

Appendix B: Results from Bid-Ask Sim	ulation
--------------------------------------	---------

CALL OPTIONS			PUT OPTIONS		
Strike price	Bid price	Ask price	Strike price	Bid price	Ask price
690	19.25	22.25	690	4	5.5
700	14	15.5	700	6.25	8.25
710	8.25	9.25	710	10	13
720	5.5	5.75	720	16.25	18.75
730	1.8	3.5	730	24	26.5
			740	32.5	36.75

Table 2: Closing bid and ask quotations at August 27, 2001, used as input in bid-ask simulation

CALL OPTIONS			\mathbf{PU}	T OPTION	IS	
Strike price	Bid price	Ask price		Strike price	Bid price	Ask price
520	37	41.25		510	7.75	10.75
530	30	33.5		520	10	12.5
540	23.75	27.25		530	12.75	15
550	18.25	21.75		540	16	18
560	14	16.5		550	20.5	23.75
570	9.25	12.5		560	25.5	29
620	0.6	4				

Table 3: Closing bid and ask quotations at September 24, 2001, used as input in bid-ask simulation

CALL OPTIONS			PUT OPTIONS			
Strike price	Bid price	Ask price	Strike price	Bid price	Ask price	
530	32.5	37.5	540	10	11.5	
540	25.75	28.25	550	12.5	15	
550	19.25	21.25	560	17.25	19.75	
560	13.25	15.75	570	22.75	25	
570	9	10.5	580	29	32	
580	5.5	7.5	590	35.75	40.75	
600	2.2	4				

Table 4: Closing bid and ask quotations at October 19, 2001, used as input in bid-ask simulation



Figure 7: Implied distribution extracted by SPLINE at August 27, 2001, together with the corresponding confidence band



Figure 8: Implied distribution extracted by DLN at August 27, 2001, together with the corresponding confidence band



Figure 9: Implied distribution extracted by SPLINE at September 24, 2001, together with the corresponding confidence band



Figure 10: Implied distribution extracted by DLN at September 24, 2001, together with the corresponding confidence band



Figure 11: Implied distribution extracted by SPLINE at October 19, 2001, together with the corresponding confidence band



Figure 12: Implied distribution extracted by DLN at October 19, 2001, together with the corresponding confidence band



Figure 13: Confidence band for SPLINE at August 27, and implied distribution extracted at September 24



Figure 14: Confidence band for DLN at August 27, and implied distribution extracted at September 24



Figure 15: Confidence band for SPLINE at September 24, and implied distribution extracted at October 19



Figure 16: Confidence band for DLN at September 24, and implied distribution extracted at October 19



Figure 17: Confidence band for SPLINE at October 19, and implied distribution extracted at November 23 $\,$



Figure 18: Confidence band for DLN at October 19, and implied distribution extracted at November 23





Figure 19: Implied standard deviation vs OBX index



Figure 20: Implied risk-neutral probability of a fall in the OBX index of 7.5 percent or more relative to the futures price the next four weeks vs OBX index (equal to Figure 5.9 in Section 5.5.1)



Figure 21: Uncertainty-parameter vs OBX index



Figure 22: Skewness-parameter vs OBX index (equal to Figure 5.10 in Section 5.5.1)

	Standard deviation	Probability-7.5% fall*	Uncertainty-parameter
Standard deviation	1	0.9816	0.9887
Probability-7.5% fall	0.9816	1	0.9798
Uncertainty-parameter	0.9887	0.9798	1
51			

*Implied risk-neutral probability of a fall in the OBX index of more than 7.5 percent relative to the futures price.

Table 5: Correlation coefficients between standard deviation, implied probability and uncertainty-parameter for SPLINE

	Standard deviation	Probability-7.5% fall*	Uncertainty-parameter
Standard deviation	1	0.8445	0.9150
Probability-7.5% fall	0.8445	1	0.9682
Uncertainty-parameter	0.9150	0.9682	1

*Implied risk-neutral probability of a fall in the OBX index of more than 7.5 percent relative to the futures price.

Table 6: Correlation coefficients between standard deviation, implied probability anduncertainty-parameter for DLN


Figure 23: Scatter plots of realized volatility versus predicted volatility (using SPLINE)



Figure 24: Scatter plots of realized volatility versus predicted volatility (using DLN, equal to Figure 5.12 in Section 5.5.2)



Figure 25: Scatter plots of realized volatility versus predicted volatility (using SLN)

	Skewness	Pearson median	Skewness- parameter
		skewness	parameter
SPLINE*:			
Upward and downward movement $(n=73)$	35~(0.592)	40 (0.175)	38(0.320)
Upward movement (n=43)	20 (0.620)	23 (0.271)	20 (0.620)
Downward movement $(n=30)$	15(0.428)	17(0.181)	18(0.100)
DLN:			
Upward and downward movement $(n=73)$	$31 \ (0.879)$	37~(0.408)	39(0.241)
Upward movement (n=43)	16(0.937)	22~(0.380)	24 (0.180)
Downward movement $(n=30)$	15(0.428)	15(0.428)	15(0.428)

The entries are the number of correct predictions in the sign test presented in Section 4.6.2. "Normal" level for all the measures (Q) are defined as the mean value. The numbers in parentheses are the significance levels from one-sided binomial tests with equal probability of success and failure.

*Using a linear extrapolation of the spline function outside the observed range of deltas (SPLINE-2).

Table 7: Number of correct predictions in sign test (significance values in parentheses). "Normal" level defined as the mean value

	Skewness	Pearson	Skewness-
		median	parameter
		skewness	
SPLINE*:			
Upward and downward movement $(n=73)$	34~(0.680)	39(0.241)	38~(0.320)
Upward movement (n=43)	20 (0.620)	23 (0.271)	$22 \ (0.380)$
Downward movement $(n=30)$	$14 \ (0.572)$	16(0.292)	16(0.292)
DLN:			
Upward and downward movement $(n=73)$	$31 \ (0.879)$	37~(0.408)	38(0.320)
Upward movement (n=43)	19(0.729)	22 (0.380)	22~(0.380)
Downward movement $(n=30)$	$12 \ (0.819)$	15(0.428)	16(0.292)

The entries are the number of correct predictions in the sign test presented in Section 4.6.2. "Normal" level for all the measures (Q) are defined as the median value. The numbers in parentheses are the significance levels from one-sided binomial tests with equal probability of success and failure.

*Using a linear extrapolation of the spline function outside the observed range of deltas (SPLINE-2).

Table 8: Number of correct predictions in sign test (significance values in parentheses). "Normal" level defined as the median value

WORKING PAPERS (ANO) FROM NORGES BANK 2001- 2002

Working Papers were previously issued as Arbeidsnotater from Norges Bank, see Norges Bank's website http://www.norges-bank.no

2001/1	Qvigstad, Jan Fredrik Monetary policy in real time Monetary Policy Wing 2001, 22p
2001/2	Claussen, Carl Andreas and Karsten Staehr Explaining the low US inflation – concidence or "new economy" Evidence based on a wage-price spiral International Department 2001, 26p
2001/3	Ball, Laurence Policy Rules and External Shocks Monetary Policy Department 2000, 21p
2001/4	Batini, Nicoletta, Richard Harrison and Stephen P. Millard Monetary Policy Rules for an Open Economy Monetary Policy Department 2000
2001/5	Guðmundsson Már, Thórarinn G. Pétursson and Arnór Sighvatsson Optimal Exchange Rate Policy: The Case of Iceland Monetary Policy Department 2001, 60p
2001/6	Leitemo Kai and Ulf Søderstrøm Simple monetary policy rules and exchange rate uncertainty Research Department 2001, 41p
2001/7	Akram, Qaisar Farooq and Ragnar Nymoen Employment behaviour in slack and tight labour markets Research Department 2001, 34p
2001/8	Kim, Moshe, Eirik Gaard Kristiansen and Bent Vale Endogenous product differentiation in credit markets: what do borrowers pay for? Research Department 2001, 24p
2001/9	Kolsrud, Dag Simulating forward-looking models Research Department 2001, 20p
2001/10	Bernhardsen, Eivind A Model of Bankruptcy Prediction Financial Analysis and Structure Department Research Department 2001, 47p

2002/1	Bache, Ida Wolden Empirical Modelling of Norwegian Import Prices Research Department 2002, 44p
2002/2	Bårdsen, Gunnar og Ragnar Nymoen Rente og inflasjon Forskningsavdelingen 2002, 24s
2002/3	Rakkestad, Ketil Johan Estimering av indikatorer for volatilitet Avdeling for Verdipapirer og internasjonal finans Norges Bank 33s
2002/4	Akram, Qaisar Farooq PPP in the medium run despite oil shocks: The case of Norway Research Department 2002, 34p
2002/5	Bårdsen, Gunnar, Eilev S. Jansen og Ragnar Nymoen Testing the New Keynesian Phillips curve Research Department 2002, 38p
2002/6	Lindquist, Kjersti-Gro The Effect of New Technology in Payment Services on Banks'Intermediation Research Department 2002, 28p
2002/7	Sparrman, Victoria Kan pengepolitikken påvirke koordineringsgraden i lønnsdannelsen? En empirisk analyse. Forskningsavdelingen 2002, 44s
2002/8	Holden, Steinar The costs of price stability - downward nominal wage rigidity in Europe Research Department 2002, 43p
2002/9	Leitemo, Kai og Ingunn Lønning Simple Monetary Policymaking without the Output Gap Research Department 2002, 29p
2002/10	Leitemo, Kai Inflation Targeting Rules: History-Dependent or Forward-Lookig? Research Department 2002, 12p
2002/11	Claussen, Carl Andreas Persistent inefficient redistribution International Department 2002, 19p

2002/12 Næs, Randi og Johannes A. Skjeltorp Equity Trading by Institutional Investors: Evidence on Order Submission Strategies Research Department 2002, 51p

2002/13 Syrdal, Stig Arild A Study of Implied Risk-Neutral Density Functions in the Norwegian Option Market Securities Market and International Finance Department 2002, 104p

KEYWORDS:

Implied risk-neutral density functions Option pricing Market expectations