


CYRIEL DE JONG

Dealing with Derivatives:

**Studies on the role, informational
content and pricing of financial
derivatives**



Dealing with Derivatives:

Studies on the role, informational content and pricing of financial derivatives

Over derivaten:

Studies naar de rol, informatieve waarde en waardering van financiële
derivaten

Proefschrift

ter verkrijging van de graad van doctor aan de
Erasmus Universiteit Rotterdam
op gezag van de Rector Magnificus
Prof.dr.ir. J.H. van Bommel
en volgens besluit van het College voor Promoties.

De openbare verdediging zal plaatsvinden op
donderdag 19 juni 2003 om 13:30 uur

door

Cyriel de Jong
geboren te Valkenburg aan de Geul

Promotiecommissie

Promotor: Prof.dr. C.G. Koedijk

Overige leden: Prof.dr. C.G. de Vries
Prof.dr. P.H.B.F. Franses
Prof.dr. F.C. Palm

Erasmus Research Institute of Management (ERIM)
Erasmus University Rotterdam
Internet: <http://www.irim.eur.nl>
ERIM Electronic Series Portal: <http://hdl.handle.net/1765/1>

ERIM Ph.D. Series Research in Management 23

ISBN 90 – 5892 – 043 –7
© 2003, Cyriel de Jong

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the author.

Gwen, voor jou

Voorwoord

Dit proefschrift is het resultaat van vier jaar onderzoek bij de vakgroep Financieel Management van de Erasmus Universiteit in Rotterdam. Het begon allemaal in Maastricht, waar ik als student door Kees werd overgehaald om met hem mee te gaan naar Rotterdam. Hoewel ik daarmee een geweldige stad en omgeving achter me liet, heb ik nooit spijt gehad van deze keuze. Het leven als onderzoeker was inspirerend, en bood mij de mogelijkheid om veel te leren, nieuwe ideeën op te doen en eigen interesses verder uit te diepen.

Vanaf het begin van mijn tijd in Rotterdam was Kees de natuurlijke begeleider van mijn proefschrift. Hoewel hij het altijd druk had, heeft hij me steeds op belangrijke momenten van waardevol advies kunnen dienen. Speciaal waren voor mij de vrijheid die hij gaf om in het buitenland conferenties te bezoeken en de nodige buiten-universitaire activiteiten te ontplooiën. Voor het vertrouwen dat hij me daarmee schonk, ben ik Kees zeer erkentelijk.

Een speciale band heb ik in de afgelopen vier jaar daarnaast opgebouwd met Chuck. Als co-auteur van het experimenten-onderzoek hebben we gezamenlijk met Kees vele ups en downs meegemaakt, maar met een resultaat dat ik beschouw als het beste van dit proefschrift. Chuck heeft me daarnaast veel bijgebracht over het uitvoeren van wetenschappelijk onderzoek, en het geduld en de precisie die daarbij horen. Op het gezamenlijke werk en mijn twee perioden in Tucson, Arizona, kijk ik dan ook met een goed gevoel terug. Thanks Chuck!

De vakgroep Financieel Management is gedurende mijn onderzoeksperiode sterk van samenstelling veranderd. Veel van die oude en nieuwe collega's hebben mijn verblijf aan de universiteit tot een heel leuke periode gemaakt. Wat na het eerste half jaar echter niet meer veranderd is, is een altijd positieve en hulpvaardige Ben als kamergenoot. Eveneens een collega, maar dan van de universiteit Maastricht, is Thorsten. Zijn enthousiasme en bedrevenheid in Gauss-programmeren hebben geresulteerd in een mooi hoofdstuk in onze beider proefschriften.

Een dankwoord ben ik daarnaast verschuldigd aan alle studenten die hebben meegedaan aan de experimenten. Hun geduld als de software het weer eens begaf, is niet voor niets geweest.

Mijn meeste dank gaat uit naar familie en vrienden, die zorgen voor de noodzakelijke ondersteuning, maar ook relativering van mijn werk. Papa, mama, Monique en vooral natuurlijk Gwen: dit boekje is er dankzij jullie.

Cyriel

Gouda, april 2003

Contents

Voorwoord (Foreword)

Chapter 1: Introduction **1**

Part I: Microstructure studies in derivatives markets

Chapter 2: Introduction to the first part **11**

Chapter 3: Stock market quality in the presence of a traded option **15**

- 3.1 Experimental design and procedures 18
- 3.2 Results 24
- 3.3 Discussion 44
- 3.4 Conclusion 45

Chapter 4: Insider strategies with options **47**

- 4.1 The model 48
- 4.2 Market quality criteria 53
- 4.3 Results 56
- 4.4 Conclusion 70

Chapter 5: Conclusion of the first part **73**

Part II: Empirical studies in derivatives markets

Chapter 6: Introduction to the second part	79
Chapter 7: The skewed-t implied distribution model	83
7.1 Methodology	86
7.2 Empirical results	91
7.3 Concluding remarks	102
Chapter 8: Implied GARCH volatility forecasting	105
8.1 Methodology	108
8.2 Data	114
8.3 Empirical results	119
8.4 Concluding remarks	128
Chapter 9: Pricing the spikes in power options	131
9.1 The two regimes model for spot electricity prices	134
9.2 Model estimation results	140
9.3 Option valuation	148
9.4 Concluding remarks	161
Chapter 10: Conclusion of the second part	165
Chapter 11: Summary and concluding remarks	169
11.1 Summary first part	169
11.2 Summary second part	170
11.3 Concluding remarks and future research	172
References	175
Samenvatting (Summary in Dutch)	189
Curriculum vitae	195

1 Introduction

Financial derivatives are the subject of this thesis. Although ‘derivative’ is a very common term in financial markets, it is instructive to dive a little deeper into its meaning. A derivative is a financial instrument whose value depends on the value of another, more basic or primitive, underlying instrument. ‘Derivative’ is in fact a very general term that can be applied to a whole range of underlying instruments and payoff structures, ranging from forwards, futures, swaps, call and put options, to complex exotic structures such as swaptions, caps, floors, straddles, spreads, butterflies and barriers. This list is constantly growing, and only limited by the fantasy of the financial community. A list of possible underlying values is equally endless and may include stocks, bonds, foreign currencies, gold, oil, electricity, credit, emission rights, transportation costs, and even weather. Why derivatives encompass such a wide variety of instruments becomes clear when one realizes that a popular underlying, a stock, may in itself already be considered as a derivative, namely as an option on the assets or profits of the issuing firm. So although a derivative seems well defined, in practice the distinction between primitive (underlying) securities and derivative securities is more diffuse. In the most general definition, every financial instrument may be termed a derivative, since its value depends on one or more factors.

Luckily, in individual cases it is often clear what is the underlying asset and what is the derivative. For example, if we price an option on a stock, then we consider the option to be the derivative and the stock to be the underlying security. Less clear however is the interaction between the two, which is the central theme of this thesis. Moving forward through this thesis we will increasingly narrow our focus. Part I (Chapters 2-5) starts with a general theoretical and experimental study on how standard options affect trading and efficiency in the underlying asset. In part II (Chapters 6-10) we first explore two specific econometric methods to infer information about future price movements in stock indices from market option prices. The final subject of part II is the most classical, in the sense that it is about derivative pricing. It presents an approach to price options on electricity spot prices, a very peculiar and risky underlying asset.

The part in this thesis on electricity options highlights the widespread acceptance of derivatives in different segments of the economy. Derivatives are used for two main and partly opposing reasons: hedging and speculation. Both are indispensable for a sound functioning market. Hedging is the reduction of financial exposure, and trade in derivative contracts offer a way to achieve this. Derivatives may facilitate risk reduction, because they are in general easier to trade and more flexible in their pay-off structures than the underlying security itself. A well-known example in which derivatives are used to hedge price risks, is an airline company with an exposure to fuel-price increases. The airline company is unable to pass fuel price rises immediately through in the flight tariffs, without losing part of the customer base. Since fuel costs represent a major component of an airline's costs, fuel price fluctuations may seriously put profitability and company viability at stake. The positive aspect of this risk is that the airline is not the only company with an exposure to it, and a very liquid market in fuel derivatives has emerged. The airline company can fixate (part of) its fuel costs by buying fuel futures or forwards, or create an upper bound on costs by buying call options. These strategies act like an insurance policy. Without a market for fuel derivatives, the only way to insure against fuel price rises would be to maintain a large stock of fuel. In comparison, derivatives provide a more convenient and cheaper solution.

The other purpose of derivative trading is speculation. While trading derivatives is only a secondary business for hedgers, it is the primary business for speculators. Speculators try to predict market movements and make a profit out of it. Unfortunately, speculation has a very negative image and is often associated with casino-like transactions. However, speculators provide the necessary liquidity for hedgers to trade and thus maintain a market. Furthermore, wise speculators carefully manage their positions to limit exposures. A well-known speculator for example is George Soros, who reportedly amassed a fortune of around €7 billion with his speculative Quantum Fund. His finest hour was on a single day in 1992, when he earned around \$1 billion by betting correctly that the British pound would fall in value. By hedging, traders bring risks to the market, whereas by speculation they bring information and liquidity to the market.

The two examples of hedging and speculation may falsely lead to the conclusion that only large investors and corporations have to do with derivatives. However, nearly every man and woman in the western world is exposed to derivatives. Money tied up in life insurance or pension funds is often managed with

derivatives. Most mortgages contain option-like provisions, whose value depends on market interest rates. And some individuals trade derivatives directly. Especially the Amsterdam options exchange (now Euronext) has traditionally been a popular venue for individuals.

Over the years, derivatives have attracted increasing attention by traders, journalists and researchers alike. Trading volumes have reached astonishing levels: on the two main European exchanges for example, Euronext – Liffe and Eurex, in 2001 a total of 216 and 675 million contracts changed hands. On Liffe's electronic trading platform derivatives trading represented a total underlying value of \$138 trillion in 2001 (figure for Eurex not published). This is around 300 times Dutch Gross National Product and growing at a much faster rate than GNP.

Instrumental to the enormous trading volumes are the advances in trading platforms. To date nearly all trades are closed electronically, either through organized exchanges or over-the-counter by derivative dealers. This explains the fast internationalization of the trading flows and the difficulty to control cross-border money flows, which is often cited as a reason for the increase in financial market volatility. The perceived volatility increase is not the only negative aspect associated with derivatives. Derivatives caused several large financial disasters and bankruptcies that received massive attention in the press. Probably most well known are the misfortunes that Nick Leeson brought to the Barings Bank (and himself ultimately) in 1995. From Barings' Singapore office he used financial futures to speculate on an increase in the Japanese Nikkei 225 stock index. Doubling his bet after each loss¹ he destroyed a total value of \$1.3 bln and induced the collapse of an old reputable London investment bank, whose clients included the British Royal family. Even more ironic for the derivative community however was the disaster concerning the hedge fund Long Term Capital Management amidst the financial turmoil of 1998. The fund had been successfully exploiting market inefficiencies for a couple of years with the help of professors Merton and Scholes, the two most acclaimed researchers in financial derivatives. Together with Fischer Black, the two professors were the bright heads behind the famous Black-Scholes option price formula for which they received the Nobel Prize for Economics in 1997. Even their brainpower could however not avoid losses accumulating to \$3.7 bln. The fund only survived thanks to a savings

¹ Brown, S.J. and O.W. Steenbeek, 2001, "Doubling: Nick Leeson's trading strategy", Pacific Basin Finance Journal, 9, p. 83-99

operation led by the US Central Bank. Confidence in derivatives in general and the professionals who use them in particular, had however been severely damaged.

The famous pricing framework that earned Scholes and Merton the Nobel Prize, relied on the assumption that derivatives are redundant assets. This means that a derivative contract can be exactly replicated with a dynamic portfolio in the underlying assets. For example, it is assumed in the Black-Scholes-Merton framework that a standard European call or put option on a stock can be replicated with a continuously updated portfolio of the stock. This assumption is convenient for pricing purposes, since its consequence is that the risk of a derivative can be hedged away and is thus irrelevant. This assumption is however at odds with all the positive and negative properties of derivatives we just discussed. If a derivative is a redundant asset, why would we trade it, and how could it cause financial disasters? The assumption of redundancy is of course a simplification; in practice, derivatives can hardly ever be replicated completely. Transaction costs, the lack of continuous trading opportunities, non-normal return distributions and non-storability of the underlying asset prohibit traders from replicating derivative trade-offs exactly. That's how derivative markets continue to flourish. They attract other traders (or the same traders at different moments) than markets in the underlying. Since both markets are interrelated, information flows from one market to the other, as we analyze in this thesis.

Part I consists of two market microstructure studies in which we explore how derivatives trading affects the price process of the underlying. More specifically, in a world of asymmetric information (some traders know more than others) we investigate how a derivative changes the quality of a market. Apart from studying information flows in general, we test whether option trading increases bid-ask spreads (trading costs), price volatility, losses to uninformed traders, and pricing errors. We use two research methodologies: experiments (Chapter 3) and a theoretical model (Chapter 4). Although there are some related experiments, the use of it is really new to this research topic. We use experiments, because they permit control of a number of confounding factors that plague empirical studies. When we study for example assets on which options are traded and compare them to assets with no related derivatives, we cannot simply ascribe differences in market quality to the trade in derivatives. Assets with derivatives are in general the more liquid assets not because of these derivatives; rather, they have derivatives because they are the most liquid and interesting to trade in. Even so plagued by natural biases are event studies that

investigate what happens to an underlying security when derivatives are introduced. The main problem with empirical studies is namely that the introduction of derivatives is not a random event, but may be biased towards more volatile assets (Vijh (1990)). To avoid such biases we use experiments where students trade in a simulated computerized market environment. This allows us to set the conditions for trading ourselves, and place all actions under a microscope. The study reveals important information flows from one market to the other, and overall indicates that an option improves the quality of a market.

Experiments and empirical studies are obsolete if theory would provide clear guidance. There are however only a few theoretical studies on this topic, and the few studies have limited applicability to real world markets. For example, two studies (Easley, O'Hara and Srinivas (1998) and John, Koticha, Narayanan and Subrahmanyam (2000)) use a single-trade setting, and thus ignore any dynamic strategic behavior or learning effects. In the theoretical model of Chapter 4 we show the importance of such dynamics. We furthermore show the mutually interacting, non-linear and sometimes opposing effects of three main variables: the number of informed trades, the proportion of liquidity motivated option trades, and the effective leverage that the option market provides to the informed trader. Though abstracting from real human behavior, the model helps to clarify what conditions lead to improvements in market quality. Just as the experiments, the model makes clear that an option market improves informational efficiency in the underlying in terms of reduced price errors (market prices are closer to their theoretical value). In an initial phase of derivative market development, liquidity motivated traders are however better off without derivatives. In the course of its development, this effect gradually reverses, but stock market volatility increases.

The first part of this thesis addresses a very fundamental issue of derivatives: are they good or bad? The second part has a slightly narrower defined topic. It presents three different option pricing methodologies that deviate from the standard assumption of normally distributed returns. In the first two main chapters, we explore two econometric methodologies to infer from option prices information about future price movements in the underlying stock index. The interest in this topic arose from the observed skews, smiles and term-structure effects in implied volatility curves. One of the most plausible explanations is that returns do not follow a Brownian motion, as assumed in the standard Black and Scholes pricing formula. The first method we explore, translates the smile and skew patterns into a skewed version of the well-

known Student-t distribution. The main advantages of this distribution are that it nests the normal distribution and that it contains direct parameters for the first three distributional moments (apart from the mean): variance, skewness and kurtosis. We describe how the parameters can be obtained and compare it with three other methods on FTSE index options. The comparison indicates that the skewed Student-t method yields a good fit to option prices, somewhat better than a popular implied tree approach. Compared to an approach that directly fits the implied volatility curve, it performs however worse. We explain this by the increased attention of option traders on implied volatility curves.

The approach that we present in Chapter 8 is more complete than the one in Chapter 7. It does not just yield a distribution of returns at one future point in time, but a complete price process. It uses the whole implied volatility term structure: options with different maturities and different strikes. We use a GARCH-type price process, for which option-pricing procedures have recently become available. The strength of the implied GARCH approach is its ability to map a risk-neutral distribution to an actual distribution, whereas the skewed Student-t approach just yields risk-neutral distribution and is thus in fact only suitable to price other options. The implied GARCH approach on the other hand can be applied to the real world. It provides traders, risk managers and market regulators with an approach to infer the market's view on actual future price movements.

The final study focuses on the energy market. Electricity markets worldwide, including all European Union countries, are in the process of liberalization and deregulation. This has opened a whole new field for financial researchers. Chapter 9 presents a methodology to price European-style options on a very peculiar underlying asset: electricity spot prices. Electricity is a pure flow variable that can only be stored at high costs. That's why its spot price² in liberalized markets is truly the result of supply and demand at that particular point in time, since no storage arbitrage can be applied. In combination with the relative inelasticity of both supply and demand, this results in prices with exceptionally high volatility, due to daily price changes of over 1000%. A few so-called 'spikes', prices that temporarily deviate largely from ordinary levels, account for a large part of this volatility. Spikes represent a non-negligible risk in electricity markets, and are extremely relevant for option pricing, but it has proven to be no sinecure to model them properly. Instead of using the popular jump-diffusion

² Spot electricity is electricity with a very short delivery period, usually one day ahead.

models, we rely on a regime-switching model. In this model the spikes represent a regime that is separate from the normal mean-reverting process. This yields a better representation of electricity price behavior and permits the separation of an option price in a mean-reverting and a spike component. Based on Dutch electricity prices, we show that the spike component may represent nearly all of the value of out-of-the-money call options. The pricing framework is relevant for the energy industry, in which this kind of options are increasingly traded to manage uncertain demand. It can furthermore be applied to price the large number of capped end-user power contracts and for the valuation of real assets, such as (flexible) power plants.

Part I

Microstructure studies in derivatives markets

2 Introduction to the first part

In this part of the thesis we investigate the implications of asymmetric information for informational linkages between an asset market and a call option on that asset. This is an interesting issue because of the empirical evidence that option listings are associated with higher market quality for the underlying asset. A plausible explanation for this effect is that the presence of a correlated asset permits the sharing of effective price discovery across markets. Market makers in the stock can set more accurate prices if they learn from transactions in the option. We use two research methodologies to study this hypothesis: experiments and a theoretical model. With the experiments (Chapter 3) we place a market in which we observe all information sets under a microscope. This allows us to examine the implications of the strategic interactions of an insider, demanders of liquidity, and suppliers of liquidity for the linkages between the two markets and the time series of prices in general. With the theoretical model (Chapter 4) we analyze on a very detailed level how the price discovery process on stock and option market interact, and how this impacts the quality of a market.

A number of insider-trading cases involving options suggest an increasing importance of options markets as an outlet for information based-trade³. At the same time, we observe a growth in the number of multiple listings of options across exchanges and an increased computerization of options trading⁴. These developments, along with increased liquidity, are accompanied by a reduction in bid-ask spreads, which make options an increasingly attractive product for informed trading.

The empirical evidence for informational linkages between the market for listed options and the underlying stock is however mixed. Fleming, Ost diek and

³ See for example “Where Have the Insider Traders Gone? Options Markets are Their New Home,” The Wall Street Journal, April 23, 1997, p. C1.

⁴ See for example “Handel op optiebeurs verstomd” (trade on options exchange has fallen silent), Volkskrant, August 27, 2002 that describes the transfer on the Amsterdam options exchange from floor-based trading to screen-based trading.

Whaley (1996) find that stocks lead options, which the authors attribute to the overall lower trading costs in the stock market. Likewise, Vijh (1990) finds that very large option trades have limited effects on option prices, and concludes that they are unlikely to be information based. Sheikh and Ronn (1994) on the other hand find systematic patterns in option returns that are inconsistent with the view of options as redundant assets. They attribute these patterns to strategic behavior by informed and discretionary liquidity traders. Similarly, Easley, O'Hara, and Srinivas (1998) find evidence that option markets are a venue for information based trading. Their results show that properly defined "bullish" and "bearish" option market volumes have predictive power for price movements in the underlying asset.

Earnings announcements are typical events around which information becomes public. Jennings and Starks (1986) and Amin and Lee (1997) study the effect of options trading on price discovery around such earnings announcements. They provide evidence that the stock price adjustment to earnings announcements is faster for firms with traded options.

There is a large body of empirical work that examines the influence of stock option listings on the time series properties of the market for the underlying asset. Much of this evidence suggests that the presence of listed options is associated with higher market quality in the market for the underlying stock. For example, Kumar, Sarin, and Shastri (1998) find that bid-ask spreads decline while quoted depth and informational efficiency increase subsequent to the listing of options. Numerous event studies find that option listing causes a decrease in volatility, although in several other studies the results are mixed or insignificant⁵. In recent subperiods Mayhew and Mihov (2000) even document increased volatility after option listings.

As shown above, the amount of empirical work on the impact of option trading is huge. Many of those empirical studies have certainly been fuelled by the lack of clear guidance that theoretical models provide. For example, Biais and Hillion (1994) show that with asymmetric information and incomplete markets, the introduction of a nonredundant option has ambiguous consequences for informational efficiency. Although the option can help avoid a market breakdown, it enlarges the set

⁵ The following studies find a decrease in volatility: CBOE (1975 and 1976), Trennepohl and Dukes (1979), Skinner (1989), Conrad (1989), Detemple and Jorion (1990), Damodaran and Lim (1991), Kumar, Sarin and Shastri (1998). In the following the results are mixed or insignificant: Klemkosky and Maness (1980), Whiteside, Dukes and Dunne (1983), Fedenia and Grammatikos (1992), Fleming and Ostdiek (1999).

of trading strategies the insider can follow, and can reduce informational efficiency by making it more difficult for market makers to interpret the information content of trades. Back (1993) extends the Kyle (1985) framework to include a call option. His main result is that the listing of an option leaves average volatility unchanged.

Of the comparatively small number of theoretical papers, at least two use a sequential trade model in the spirit of Glosten and Milgrom (1985): Easley, O'Hara and Srinivas (1998), and John, Koticha, Narayanan and Subrahmanyam (2000). Both models are developed in an asymmetric information setting in which informed traders may trade in stock or option markets. Easley et al. (1998) study whether option markets attract informed trading and whether they incorporate information more quickly than stock markets. They derive that under certain conditions options are attractive to traders with superior information. John, Koticha, Narayanan and Subrahmanyam (2000) use a single-trade model to study the impact of option trading and margin rules on the microstructure of stock and option markets. They analyze opening quotes and show that the introduction of option trading increases quoted spreads, but improves the informational efficiency of stock prices irrespective of whether or not binding margin requirements are in place. The increase in the informativeness of the trading process results because with option trading private information can be inferred from two sources – order flow in the stock and option markets.

In the next chapter we use a controlled economic experiment where students trade in a laboratory environment. In that chapter we study the implications of derivative trading in an asymmetric information setting. The use of laboratory asset markets is relatively new to the market microstructure research, and the derivative research in particular. Although experimental research poses some difficulties (see for example Kagel and Roth (1995), and Friedman and Sunder (1994)), we will show in the next chapter that it may provide valuable insights. With experiments, the researcher can change some aspects of the economic environment while keeping other aspects constant. The researcher is furthermore able to control the amount of each trader's information. Finally, experiments permit the researcher to study detailed trading information, not only including executed trades, but also all dealer quote revisions, and profits and losses incurred by the participants in the market.

With the experiments we present a unique approach to study the overall effect of option trading on market quality in the underlying. In order to better understand the different mechanisms that lead to the results, in Chapter 4 we set up a dynamic

theoretical model that resembles the experimental design. The model is sequential, which means that only one trade is executed at a time, and has a similar design as the single-period models in Easley, O'Hara and Srinivas (1999) and John, Koticha, Narayanan and Subrahmanyam (2000). The experiments made convincingly clear that the effect of options must be analyzed in a dynamic setting, and not in a single-trade setting as do those theoretical studies. Options allow informed traders to strategically switch between two markets: traders are for example willing to forego immediate profits for larger profits at a later point in time. At the same time, dealers can infer information from both markets to set more accurate quotes. We therefore extend the single-trade model into a multi-trade model in which traders may trade more than once in a period and dealers update information after each trade. Because of the increased model complexity, outcomes are not available in closed form, but generated by simulating trading sequences.

With the experiments we obtain insights beyond empirical studies in a relatively realistic setting including real human behavior. The theoretical model on the other hand is more stylized, but yields additional insights. In the model we assume for example that dealers are completely rational, competitive and risk-neutral, whereas in the experiments risk-aversion, fierce competition and collusion between dealers play a role. Abstracting from these behavioral aspects, the simulations provide new insights in the complex price discovery process that determines the quality of a market.

3 Stock market quality in the presence of a traded option ⁶

Theoretical evidence on the implications of options trading for market quality of the underlying does not provide clear guidance, and the empirical evidence for informational linkages between both markets is mixed and may be biased. Vijh (1990) for example makes an important note on all the event study results around option listings: because the listing of options is by no means a random event, the research results may be biased, especially when options are most likely to be introduced on options and in periods with high volatility. To circumvent the problem of empirical studies to control for such biases, we use a controlled economic experiment. The experimental results clarify the implications of asymmetric information for informational linkages between a stock market and a traded call option on that stock.

Our experimental design incorporates both an asset and a call option on that asset (for simplicity we refer to the asset as a stock). The framework is based on the Kyle model with a single insider who knows the ex post liquidation value of the asset. Liquidity shocks in the stock and option are exogenously determined as in Back (1993), although we permit liquidity traders to act strategically in an attempt to minimize trading costs. We also employ a transparent quote-driven trading protocol, with all trades executed by competing dealers, and all trades visible to all participants. We thus obtain a standard trading mechanism, and impose no constraints on the strategies of the insider, dealers or liquidity traders.

One experimental approach would be to compare a market with a traded option to a market operating in isolation. We do not attempt this since it would introduce control problems, as follows. Adding a traded option with dedicated option dealers would increase the supply of liquidity services and confound our attempt to isolate differences due to information flows. Adding a traded option without dedicated

⁶ This chapter is based on: C. de Jong, C.G. Koedijk and C.R. Schnitzlein, 2002, “Stock market quality in the presence of a traded option”, working paper.

dealers (requiring the stock dealers to make a market in both the stock and the call simultaneously) would dramatically increase the difficulty of the dealers' task.

Rather than take this approach we focus on the role of the ex post intrinsic value of the option. We set the strike price of the option equal to the (uninformed) expected value of the stock. With the stock distribution approximately Gaussian, the option distribution is highly skewed around its expected value. When the intrinsic value of the option is positive, the linkage between the two markets is direct in the sense that knowing the liquidation value in one market is perfectly informative with respect to the liquidation value in the other. The potential for option trading to contribute to price discovery in the stock here is large. When the option's intrinsic value is zero, knowledge of the option liquidation value only permits a truncation of the stock value distribution, with a corresponding decoupling of price discovery. Therefore, the informational linkage between the two markets depends on the intrinsic value of the option. This allows an examination of how price discovery in the markets depends on the presence of a highly correlated asset. We argue below that differences in market quality for the underlying asset as a function of the ex post intrinsic value of the option will provide evidence for how and why the introduction of a traded option matters.

Our major findings are the following. The insider trades aggressively in both the stock and the option, and typically trades in the market that affords the most profitable trading opportunity. Liquidity traders concentrate their trades at the end of the trading period when trading costs are low, and insider trading patterns mimic those of liquidity traders. Price discovery with respect to intrinsic value takes place in both markets and this leads to informational linkages: dealers in each market revise quotes in the direction consistent with an information story in response to trades in the other.

A strong set of results pertains to the effect of the intrinsic value of the option on liquidity and price efficiency. When the intrinsic value of the option is positive, convergence to intrinsic value in the stock is faster, liquidity trader losses in the stock are smaller, and the volatility of transactions prices is lower than when the option's intrinsic value is zero. This appears due to the greater information content of trades in the option market in this case: this information is used by the dealers in the stock to more rapidly pinpoint intrinsic value. Of particular importance here is how differences in the dealers' conditional expectations for the stock and option values interact with the strategy of the insider to increase the message space of the dealers, and help dealers more rapidly pinpoint fundamental value.

Although the experimental asset market literature is large and growing, there are relatively few studies that incorporate derivatives. For example, Forsythe, Palfrey and Plott (1982, 1984) and Friedman, Harrison, and Salmon (1984) find that in a private-value setting⁷ the inclusion of a futures market speeds convergence to informationally efficient pricing, although while the latter study finds that futures reduce spot-price variability, the former study does not. Plott and Sunder (1988) compare private-value markets organized with two or three trader types and three states with two other types of markets. In the first there is a single trader type (common values) and in the second there are three option-like contingent claim securities, each paying off in one (and only one) state. They find that the contingent claims markets aggregate information better than the single security market and propose possible explanations including the importance of knowledge of others' preferences as a necessary condition for the aggregation of diverse information. However they conclude their understanding of the issue is so incomplete they cannot even provide a precise conjecture (p. 1116). Finally, Kluger and Wyatt (1995) find that in a two-state private-value setting, an alternating asset market and option market (that are never open simultaneously) converge to equilibrium faster than an asset market alone. What distinguishes our experimental design from previous experimental work therefore is its focus on information flows in a common-value setting, where the option market and the stock market are open continuously.

A barrier to experimental research that incorporates options is the necessary complexity of such markets. In order to ensure that our subjects could master the trading environment and employ sophisticated strategies we used only advanced finance students who were already familiar with options and their associated non-linear payoffs. We also developed and employed custom software that aided subjects in mastering a setting that involves the trading of two interdependent assets. Finally,

⁷ In a private-value experimental setting, a motive for trade is induced by introducing a limited number of states (two or three) and giving two or three different trader types state contingent payoffs that are negatively correlated (e.g., trader type I receives a high payoff in state A and a low payoff in state B while type II traders have payoffs that are reversed). Typically subsets of traders are given information with respect to which states can or cannot occur in a given period. Allocative efficiency measures the extent to which the trader type with the highest state contingent valuation holds the assets (which are in fixed supply) at the end of the period, and price efficiency measures the extent to which prices converge to the highest dividend in the realized state.

we relied on extensive training and multiple replications in order to verify the effect of experience on outcomes: each cohort of subjects was kept intact for a minimum of four experimental sessions in addition to an initial training session.

The plan of the chapter is as follows. In Section 3.1 we describe the experimental design and procedures. In Section 3.2 the data are presented and analyzed. In Section 3.3 we discuss our results in the context of theoretical models and in Section 3.4 we conclude.

3.1 Experimental design and procedures

Our experimental design is roughly based on the Kyle (1985) model, generalized to include a single call option. There is a single insider that knows the intrinsic value of the stock and two liquidity traders who receive asset specific exogenous liquidity shocks in both the stock and the call. Three competing dealers maintain outstanding bids and asks for the stock and three additional dealers make a market in the call.

We impose no constraints on the timing, number, or direction of trades the insider may make in either market. For example, the insider is permitted to trade against her information in either market in an attempt to camouflage her information. We do not impose borrowing constraints or short-sales constraints, so leverage effects that might make options attractive to informed traders are not present in the experimental markets. The insider's choice to trade in the option or the stock is dictated entirely by the relative magnitude of profitable trading opportunities in the stock and the option, and the price responsiveness to order flow implied by the timing and direction of her trades. We do limit all trades to a single unit in order to reduce the difficulty of the dealer's task, and sharpen our ability to make inferences with respect to the informational content of the order flow.

As is standard in the microstructure literature (and more specifically in microstructure models that incorporate an option, e.g., Back (1993), Easley, O'Hara, and Srinivas (1998), and John, Koticha, Narayanan, and Subrahmanyam (2000)) we impose liquidity shocks in each market that are exogenous to intrinsic value. This avoids a "no trade" equilibrium as in Milgrom and Stokey (1982) and Biais and

Hillion (1994).⁸ An exogenous liquidity trading motive is particularly natural in options markets where many trades are hedging related and by their nature non-informational. We operationalize liquidity shocks by giving each liquidity trader a required end-of-period position in both the stock and the call. We make these demands inelastic by imposing penalties of sufficient magnitude that are proportional to the distance between a liquidity trader's final position and his required position. Although their required positions are exogenously determined each period, there are no constraints on how liquidity traders arrive at these positions: they are free to choose the timing of trades in both markets in order to minimize trading costs.⁹ We also impose larger liquidity shocks on average in the stock market. This matches a feature of field markets, and is a possible impediment to an equilibrium in which the insider trades the option.¹⁰

We employ a transparent quote-driven trading protocol: dealers maintain standing bids and asks in separate centralized limit-order books for the stock and the option. This is a standard feature of many field markets, and reduces the complexity of the task for subjects since they transact against known prices. A dealer is free to compete for order flow by adjusting his bid or ask at any time and there are no

⁸ The imposition of exogenous liquidity shocks has also become a standard feature of market microstructure experiments. The use of computerized liquidity traders (Schnitzlein (1996), Bloomfield and O'Hara (1998)) has the advantage of reducing the required number of subjects while allowing the perfect control of the timing and size of liquidity shocks. In this study, subjects are employed in the role of liquidity traders to allow the examination of issues pertaining to their strategic timing of trades in order to minimize transactions costs. Other studies that employ strategic liquidity traders include Lamoureux and Schnitzlein (1997), and Cason (2000). Finally Bloomfield and O'Hara (1999, 2000) employ both strategic and computerized liquidity traders.

⁹ We give each liquidity trader uncorrelated required positions in the stock and the option for experimental expediency: it reduces the number of agents required - an important consideration since our design required keeping cohorts of subjects intact for four sessions. This design feature does not change the incentives of liquidity traders to act strategically in order to minimize trading costs.

¹⁰ In Easley, O'Hara, and Srinivas (1998), whether an informed trader only uses stocks or trades both stocks and options depends in part on the fraction of liquidity traders in the option. Theory is silent on the relationship between the magnitude of liquidity shocks in the stock and option. Relative stock and option volumes suggest that non-informationally motivated trade is probably larger in absolute terms in stocks than options. We therefore match this feature of field markets with our experimental design.

constraints or costs associated with quote revisions. Unlike most theoretical models where dealers are assumed to be risk-neutral agents that set quotes equal to conditional expectations, dealers in the experimental markets are motivated to engage in price discovery in order to earn trading profits.

We include only one call option (and no put options) to create the following asymmetry. When the intrinsic value of the option is zero, knowledge of the option liquidation value only permits a truncation of the stock value distribution. When the intrinsic value of the option exceeds zero, the liquidation value in one market is perfectly informative with respect to the liquidation value in the other. This design allows us to analyze how the introduction of a correlated asset like an option affects market quality while enhancing experimental control: since the same group of subjects participates in multiple market periods when the intrinsic value of the option is either positive or zero, we control for the influence individuals might have on market quality as a function of (positive or zero) intrinsic option value.

3.1.1 Parameter values and variable distributions

Agents begin the first market period of a session with a cash balance that differs across the type of agent. Cash balances and all other values are expressed in laboratory dollars (L\$). The insider has a starting cash balance of L\$450, the stock and option dealers of L\$550, and the liquidity traders of L\$900. The differences in starting cash balances are intended to minimize differences in profits by trader type. The starting cash balances average L\$617. Trading profits (losses) are carried forward to subsequent periods. At the end of the final market period, the cash balances are multiplied by 0.08 to convert laboratory dollars to Dutch Guilders (NLG)¹¹, and each subject is privately paid his earnings. Given the zero-sum nature of the trading game, cash payments (net of any penalties incurred by traders) average L\$617 (NLG 50) per subject per session.

Each market period the stock value is known to be drawn from an approximate normal distribution with mean of L\$100, standard deviation of L\$12, and support on the whole laboratory dollars between L\$50 and L\$150, both inclusive. The

¹¹ One Dutch Guilder equals approximately € 0.45, or US\$ 0.50 at the time the experiments took place.

value of the European call option with a strike of L\$100 is $\max[\text{Stock} - 100, 0]$. A riskless asset (cash) is included that for simplicity pays an interest rate of zero.

3.1.2 Traders

Four different types of agents are present in each market period: one insider, two liquidity traders, three stock dealers and three option dealers. We define a cohort to be the nine subjects that always trade together. Every cohort participates in four or five different sessions. Each session consists of between five and nine different market periods, with each market period defined by a different set of random draws for the asset values and liquidity shocks.

The insider is the only agent to learn the end-of-period stock and option value prior to the start of the trade. Each of the two liquidity traders is required to finish the trading period with a randomly determined position in the stock and the option. Each trader's required position is determined by an independent draw from a discrete uniform distribution. The draws range from -6 to $+6$ for the stock, and from -3 to $+3$ for the option, each integer-valued and excluding 0. If a trader does not meet this requirement, a penalty is assessed at the end of the period equal to L\$100 times the absolute value of the deviation between the required position and the actual end-of-period position. The magnitude of the penalty ensures that demand is inelastic at the required position. Each period, total trading profits and losses of each trader are added to starting cash balances and carried forward to the next period. Therefore, traders have an incentive to minimize trading costs. Each liquidity trader privately learns his required position prior to the start of each market period.

Three stock dealers make a market in the stock and three option dealers make a market in the option. The dealers have to provide quotes at all times, with the only restriction that they must lie in the interval $[50, 150]$ for the stock and $[0, 50]$ for the option. Dealers see the quotes reported by all other dealers and are allowed to revise their quotes at any time. Apart from providing bids and asks, they can initiate trades themselves as well, either in the market for which they provide quotes, or in the other market. At the beginning of each market period, dealers do not learn either the end-of-period stock or option value nor the liquidity traders' required positions, and they do not receive required positions. All information pertaining to distributions, parameters, and the rules governing trade is common knowledge.

3.1.3 Communication

All interactions among subjects are conducted on a series of networked personal computers with custom software. The computer screen in front of each subject contains continuously updated market information, including trade history, cash balance, stock and option positions, net market order imbalance (buyer initiated trades less seller initiated trades), current bids and asks, and the time remaining in the current trading period. The trade history shows the trades that have occurred, whether a buyer or seller initiated them, the price of the trade, and each agent's own transactions. In addition, the insider's screen shows the end-of-period value of both assets, and each liquidity trader's screen shows his end-of-period required positions for the stock and option.

3.1.4 Trading procedures

Before the trading interval can begin, each dealer in the stock and each dealer in the option must submit a bid and an ask.¹² A dealer's bid represents the price at which he/she is willing to buy a single unit of the risky asset while a dealer's ask represents the price at which he/she is willing to sell a single unit. Only the stock dealers submit quotes in the stock market, and only the option dealers submit quotes in the option market. After all dealers have entered their quotes, the markets open and the trading interval clock begins a 60-second countdown. Each dealer does not observe the other dealers' quotes until the market opens. During the trading period traders are free to hit inside quotes in both markets, and each dealer is free to revise his outstanding quotes at any time. The quotes are displayed on each agent's computer screen so that bids are in ascending order and asks are in descending order, with the inside quotes highlighted. A dealer can move to the "inside" on either side of the spread by improving on the current inside quote. When a trader initiates a trade at the inside bid or ask, all traders and dealers observe the transaction and the price, but they do not learn which trader initiated the trade. The market period clock is stopped while the dealers go through the process of resubmitting quotes. Therefore, the actual

¹² Starting bids and asks in the stock must bracket the expected value of L\$100, starting bids and asks in the option must bracket the expected value of L\$4.78.

trading interval requires considerably more than 60 seconds: most took between ten and twelve minutes. Trades represent a zero-sum game, with the profitability of each trade determined on the basis of the relation between the trading price and the end-of-period asset value. Trading profits on a trade are equal to the signed trade (+1 for a buy, -1 for a sell) times the end-of-period asset value less the transaction price. All transactions are for a single unit of the risky asset.

At the end of each market period each subject is informed of the actual end-of-period value of the stock and the option, personal trading profits (or losses), any penalties incurred (in case of the liquidity traders), and his final cash balance after liquidation of end-of-period positions at stock and option intrinsic value.

3.1.5 Subjects and experimental procedures

The participants in the experiments are students in the Rotterdam School of Management at Erasmus University. All were students with a specialization in finance (32 undergraduates and 4 graduate students). The first set of markets involving cohort I was conducted in September of 1999, followed by a set of markets involving cohort II in January 2000. In order to verify our results, we formed cohorts III and IV and ran an additional set of markets in March and April 2001.

Each cohort went through training in the information structure, variable distributions, and parameter values, and on average five market periods of trade. The training round lasted between two and three hours, and ensures that subjects possess a good understanding of the rules of the market.

Each session was preceded by a review of trading rules, parameter values and distributions. The review became progressively shorter as the subjects became more experienced. In the first session, subjects were randomly assigned to roles. In the second to fourth sessions, the roles of subjects were changed according to a predefined scheme of which they were unaware: subjects are insiders at most once, and liquidity traders, stock dealers and option dealers at least once and at most twice. The benefit of this role-switching scheme is that it controls for differences across participants - we do not want those differences to drive differences between trader types – and it provides a

deep understanding of the market setting that comes from assuming the roles of different agent types.¹³

3.2 Results

Each of the four cohorts participated in an initial training session and then four (or in one case, five) two-hour sessions. Since subjects were assigned to new roles at the beginning of each session, we exclude the first two market periods from each session in order to allow subjects to refamiliarize themselves with their roles, the trading rules, and the software. This yields an initial data set of 17 sessions, 132 market periods and 3,475 transactions.

Our initial data analysis revealed mild learning effects before the third session: dealer losses are sometimes large over the first two sessions, with the dealers on average incurring losses. In contrast over the last two sessions the dealer industry is profitable in all four cohorts. Since dealers' quote-setting strategies are completely unconstrained (dealers may set bids and asks at the support of the asset value distribution – prices that cannot be unprofitable) dealer losses are clearly inconsistent with equilibrium behavior. Since our intention is to present “equilibrium” results, we therefore report analysis based on the last two sessions from each cohort. Our “experienced” data set consists of 8 sessions, 62 market periods, and 1,630 transactions. As a robustness check we also performed the analyses that follow with all 17 sessions. Major results are unchanged although on most dimensions behavior shows more variation when the initial sessions are included.

Each market period in each session is distinguished by a distinct set of random draws for the stock value, the option value, and the liquidity traders' demands. Summary statistics pertaining to these draws are reported in Tables 3.1 and 3.2.

¹³ Ball, et al. (1991) find that convergence to equilibrium strategies in bargaining games with adverse selection is speeded by rotating the roles of subjects.

		# Trades / Period						Profits								
Asset value		Stock			Option			Stock			Option					
Average		Liq.			Liq.			Liq.			Liq.					
Session	Periods	Stock	Option	Insider	traders	Insider	traders	Insider	traders	Insider	traders	Insider	traders	dealers	Option	dealers
1-1	7.0	100.1	3.7	8.1	7.4	7.1	3.7	44.5	-70.7	25.5	11.4	-9.5	-2.2			
1-2	7.0	99.7	4.7	4.9	9.1	6.0	3.4	26.7	-35.9	9.6	11.1	-14.6	3.6			
2-1	8.0	100.1	4.9	7.0	7.9	3.3	3.9	55.2	-64.9	8.8	17.2	-11.2	-5.1			
2-2	8.0	98.1	4.4	5.9	8.0	4.1	4.0	49.8	-66.4	16.8	22.8	-5.5	-13.5			
3-1	8.0	104.1	8.6	9.3	7.4	3.1	3.6	33.4	-53.7	16.4	6.7	-12.9	15.6			
3-2	8.0	98.6	3.4	4.1	7.8	1.8	4.4	29.7	-13.7	-14.1	3.0	-3.7	0.2			
4-1	8.0	97.0	2.6	5.9	7.5	4.8	4.0	52.8	-99.7	42.4	8.1	-9.6	3.8			
4-2	8.0	97.3	3.8	4.5	7.3	5.5	4.1	20.8	-55.0	39.5	1.0	-11.3	10.2			
Mean		99.4	4.5	6.2	7.8	4.4	3.9	39.1	-57.5	18.1	10.2	-9.8	1.6			
Std.Err. (Dev)		-11.7	(6.5)					4.6	9.0	6.4	2.5	1.3	3.2			

Table 3.1. *Asset values, trading activity and profitability.*

This table reports asset values, trading activity and profitability for each of the 62 market periods, aggregated by session.

A principal finding that we analyze in detail below is that price efficiency in the market for the stock is higher when the intrinsic value of the option is positive (in-the-money). This effect is very strong, is present in all four cohorts, and is highly significant when we use a matched-pairs t-test to compare a single mean level of price efficiency for each cohort when the ex post value of the option is in-the-money, with a single mean when the option is out-of-the-money ($t = 6.86, p < 0.01$).¹⁴ A comparably strong result obtains for the effect of a positive intrinsic option value on stock price volatility ($t = 6.48, p < 0.01$). Later in this section we analyze the behavior that leads to these main results. This requires analysis on the level of the session, market period, and in some cases individual transactions. In order to account for possible session-level interdependencies in the data, we report all regressions with p-values that are based on GMM t-statistics statistics that are consistent in the presence of heteroskedasticity and autocorrelation.¹⁵ All reported p-values are for two-tailed tests.

3.2.1 Stock market quality and the intrinsic option value

In this section we characterize stock market quality as a function of whether the ex post intrinsic value of the option is positive or zero. Our initial measures of market quality are price efficiency, volatility, and liquidity trader losses.

We define price errors (PE) as the difference between intrinsic value and the transaction price, in absolute value. We also report price errors relative to the midpoint of the inside bid-ask spread (average of highest bid and lowest ask) in absolute value. Relative price errors provide a measure of how rapidly information in the order flow is

¹⁴ We report both t-tests and results from the non-parametric randomization (permutations) test. With four cohorts, the randomization test yields the theoretical minimum p-value for a two-tailed test (0.125).

¹⁵ Analysis on the level of the cohort is most conservative since observations across sessions are by definition independent. We test for cohort effects (dependencies across sessions within cohorts) by analyzing realized bid-ask spreads. We choose bid-ask spreads because although we change the roles of agents between sessions, the task of dealers is most complex and potentially subject to the influence of previously observed behavior. We test for dependencies by performing both parametric and nonparametric ANOVA on average session realized spreads in the option, with the data grouped by cohort. None of the tests approach even marginal significance. The variation in behavior across sessions is primarily due to the agents that assume the role of dealer, and not the cohort to which the dealers pertain.

incorporated into market prices in the presence of a strategic insider who chooses between transacting in the option and the stock.

When the intrinsic value of the option is positive, mean price errors in the stock market are significantly lower. This effect is very strong, and is present at the level of cohorts, sessions, market periods, and individual transactions. We augment the cohort level analysis reported above by calculating for each session both a “within session” mean of the periods in which the intrinsic value of the option is zero, and a “within session” mean when the intrinsic value is positive. In seven out of eight sessions mean price errors are smaller when the intrinsic value of the option is positive¹⁶. Both a matched pairs t-test and the non-parametric randomization test yield a highly significant difference ($p < 0.01$).

Realized spreads in the stock (ask minus bid at the time of each transaction) are only slightly smaller when the intrinsic value of the option is positive. Aggregating on the level of the session as above, the difference is insignificant. The reduction in realized spreads does not explain the increase in price efficiency: the informational efficiency of the midpoint of the bid-ask spread also increases when the intrinsic value of the option is positive ($p = 0.03$ and $p = 0.06$ for the t-test and randomization tests respectively). Mean price errors measured relative to transaction prices and spread midpoints, and realized spreads by session are reported in Table 3.2.

¹⁶ The first session of cohort 4 is the only exception, and can be explained as follows. In all three periods in which the option is in-the-money, both liquidity traders had to sell the stock. As a result, all liquidity traders traded in the opposite direction from the insider, complicating price discovery.

Sess.	Per.	Avg. stock	Stock trades	Option trades	PE _S (price)	PE _S (mid)	PE _C (price)	PE _C (mid)	PE _{SC} (price)	PE _{SC} (mid)	Stock spread	Option Spread	DP _S (mid)
Panel A: Intrinsic value of the option is zero													
1-1	5	95.0	15.2	12.6	9.9	10.7	2.0	5.2	6.3	8.2	16.2	8.7	1.8
1-2	4	91.3	18.8	8.8	6.7	8.2	3.6	4.7	5.7	7.1	11.1	6.7	1.3
2-1	4	90.5	18.8	8.8	11.2	11.9	5.5	7.5	9.4	10.5	9.9	9.0	1.2
2-2	6	91.7	16.3	8.5	11.1	10.7	2.9	3.7	8.3	8.3	10.5	3.7	1.6
3-1	3	88.0	20.3	9.7	8.9	11.1	1.8	4.3	6.6	8.9	15.4	7.0	0.9
3-2	5	92.4	13.2	6.2	7.3	8.4	1.4	3.7	5.4	6.9	8.2	6.4	1.0
4-1	5	91.0	15.6	12.6	12.4	13.8	1.4	3.5	7.5	9.2	13.2	5.4	2.1
4-2	5	89.6	12.8	12.6	9.3	9.8	2.8	4.2	6.1	7.0	17.3	5.9	0.8
Avg		91.2	16.4	10.0	9.6	10.6	2.7	4.6	6.9	8.3	12.7	6.6	1.3
Panel B: Intrinsic value of the option is positive													
1-1	2	113.0	20.5	15.5	7.0	11.0	7.7	7.5	7.3	9.5	14.6	7.3	1.6
1-2	3	111.0	10.3	12.0	4.3	3.4	3.9	6.4	4.1	5.0	10.6	9.4	1.1
2-1	4	109.8	15.3	9.5	5.8	7.3	4.5	5.5	5.3	6.6	10.3	7.6	1.0
2-2	2	117.5	9.5	12.5	10.1	11.4	10.6	11.9	10.4	11.7	8.7	5.5	1.5
3-1	5	113.8	21.8	12.2	5.3	8.5	4.7	6.3	5.1	7.7	12.4	10.8	0.8
3-2	3	109.0	13.0	8.7	4.0	4.2	3.0	4.7	3.6	4.4	7.8	7.5	0.7
4-1	3	107.0	14.3	14.3	13.0	9.9	5.4	5.3	9.2	7.6	17.3	6.2	1.0
4-2	3	110.0	12.7	8.7	5.4	10.3	5.4	6.1	5.4	8.6	15.2	6.8	0.6
Avg		111.4	14.7	11.7	6.9	8.3	5.7	6.7	6.3	7.6	12.1	7.6	1.0
Difference			1.7	-1.7	2.7	2.3	-3.0	-2.1	0.6	0.7	0.6	-1.0	0.3
p-value: matched pairs t-test					0.4%	3.0%	2.1%	7.2%	43.5%	49.7%	45.5%	14.9%	3.7%
p-value: matched pairs rand. test					70.0%	5.8%	1.7%	3.5%	41.9%	47.8%	45.7%	15.9%	0.0%

Table 3.2. *Price efficiency, spreads, volatility and the intrinsic value of the option.*

This table reports price errors, realized bid-ask spreads and volatility for each of the eight sessions in the experiments. Panel A is based on the periods where the end-of-period intrinsic value of the option is zero; Panel B is based on the periods where the end-of-period intrinsic value is positive. PE_S (price) is the mean of the absolute difference between the stock transaction price and the intrinsic stock value. PE_S (mid) also measures price errors in the stock, but relative to the midpoint of the bid-ask spread. PE_{SC} (price) is the mean of the absolute difference between the transaction price and the stock or option intrinsic value depending on the asset traded, and is our measure of price efficiency in the aggregate. PE_{SC} (mid) is defined similarly. Stock (Option) spread is the realized stock (option) bid-ask spread at the time of a stock (option) transaction. DP_S is the mean change in the midpoint of the bid-ask spread, and is a measure of market volatility. In all cases, an average is calculated for each market period in a session, and then each market period is weighted equally. When option intrinsic value is positive, stock price errors and volatility are significantly lower, and option price errors are higher.

What causes the dramatic increase in price efficiency when the option is in the money? The short answer is that the insider's trades in the option are more informative only when the option's ex post intrinsic value is positive. This in turn allows dealers in the stock market to benefit from price discovery in the option market. We analyze this phenomenon in detail in Section 3.2.2.

Higher price efficiency when the intrinsic value of the option is positive is also associated with lower stock price volatility. Our measure of volatility is constructed as follows. First we calculate the change in the midpoint of the bid-ask spread for each transaction in a market period ($|mp_t - mp_{t-1}|$), and then compute the average for each market period. Spread midpoints are used in order to control for bid-ask bounce, with changes in the midpoint serving as a proxy for the transaction-induced revision in the conditional expectation of the intrinsic stock value.¹⁷ Market period averages are calculated for each session and reported in Table 3.2. Weighting each session equally, mean midpoint price changes are 27% less when the intrinsic value of the option is positive ($p=0.04$ and $p<0.01$ for the t-test and randomization tests respectively). This effect is present in all eight sessions (and all four cohorts), and indicates a more direct convergence to intrinsic value in the stock.

Our third measure of stock market quality is the magnitude of liquidity trader losses. Liquidity traders have inelastic demands that are uncorrelated with intrinsic value, and as expected incur large significant losses from their stock market activity. These losses decline by 44% when the intrinsic value of the option is positive (Table 3.3). This effect occurs in six of the eight sessions and is due to liquidity traders on average trading at prices closer to intrinsic value. In the two sessions where it does not occur, it is explained by the average direction of liquidity trades relative to the intrinsic value of the stock (for example, in the first session with cohort 4, in all three periods in which the option is in-the-money, both liquidity traders had to sell the stock, resulting in very large liquidity trader losses). In Table 3.4 we present an analysis that controls for the direction of liquidity trades relative to the value of the stock and option, and we find a highly significant reduction in liquidity trader losses in the stock in the presence of an option with positive intrinsic value ($p<0.01$).

¹⁷ Inventory effects may imply this conditional expectation need not be centered at the spread midpoint. We assume these effects are of second order importance.

3.2.2 Price errors and time

In this section we examine the role of the option’s intrinsic value on the convergence to informationally efficient pricing in the stock using transactions level data and a parsimonious specification. In every market period, the trading activity of the insider causes stock and option prices converge to intrinsic value by the end of the market period; our focus here is on how the intrinsic value of the option affects the speed of convergence to strong-form efficiency.

Recall that at the beginning of a market period, dealers know that the probabilities of the option being in or out-of-the-money at the end of the period are roughly equal. We regress stock price errors on a constant, the absolute difference between the stock value and its unconditional expectation (DEV), the product of the net liquidity shock in the stock and the difference between the stock value and its unconditional expectation (COR(S)),¹⁸ an indicator variable that indicates whether or not the insider initiated the trade (INSIDER), the transaction time (TIME), and the transaction time interacted with a variable that indicates whether or not the intrinsic option value is positive (TIME*ITM). To derive the transaction time we order the stock transactions and divide them by the total number of transactions in a period. This definition of time accounts for the clustering of transactions in clock time. The number of observations is 974 and the adjusted R² is 25.8%.

$$PE = b_0 + b_1DEV + b_2COR(S) + b_3INSIDER + b_4TIME + b_5TIME*ITM \quad (1)$$

10.32	0.20	-0.03	-3.51	-3.79	-4.87
(9.09)	(2.62)	(-4.57)	(-4.83)	(-2.57)	(-3.41)

All parameter estimates have the expected sign and are significant at the 1% level. When the asset value is extreme, price errors are significantly higher: in this case convergence to informationally efficient pricing requires more time. When the net liquidity trade is in the direction of the insider’s information price errors are lower. This is due to fewer trades on the “high price error” side of the spread as price converges to intrinsic value, and a simpler information extraction problem for the dealers. When the insider transacts price errors are lower, because her trades are on the “low price error” side of the spread as insider activity moves price toward intrinsic

¹⁸ This variable measures the extent to which the liquidity traders trade in the same direction as the insider’s information.

value: it is the cumulative effect of insider trading that causes price errors to decline with time.

The coefficient on the time of trade variable interacted with the positive option value indicator is negative and highly significant. This indicates more rapid convergence when the option is in-the-money. This is because when the option is in-the-money, the option value is perfectly correlated with the stock value: trades in the option help dealers pinpoint the value of the stock. When the option is out-of-the-money, trades in the option only aid the dealers in truncating the stock value distribution. We analyze in detail in Section 3.2.6. below the informational linkages between the two markets.

3.2.3 Insider profits and behavior

Market period data on profitability by trader type are reported in Table 3.1. Profits on each trade are $(V-P)*Q$, where V is the end-of-period value of the asset, P is the transaction price, and Q is an indicator (+1 for a buy, -1 for a sale). An agent's market period profits are the sum of the profits on all transactions that agent participated in during that market period. Because the insider is the only agent who knows the terminal value of the stock and option, the behavior of the insider is critical in determining patterns in informational efficiency.

The profits earned by the insider are significantly greater than zero in both the stock market and the option market. In most market periods (56%) the insider earns profits in both the stock market and the option market. Over the 62 trading periods, the insider traded 384 times in the stock and 234 times in the option.

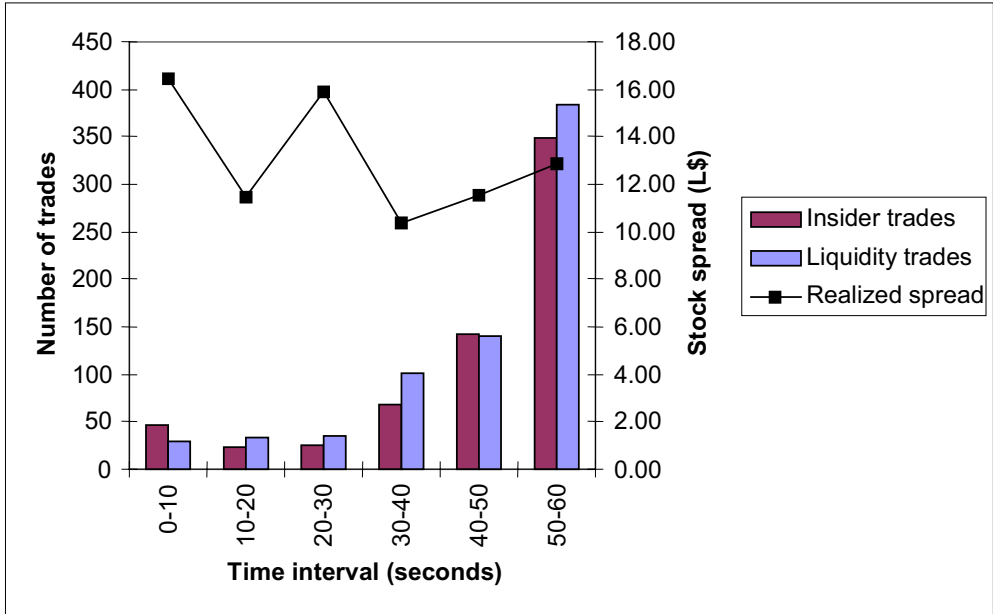


Figure 3.1. *Temporal patterns in trading activity and the bid-ask spread.*

This figure shows the concentration of both insider and liquidity trader activity at the end of the trading period and the temporal decline in realized bid-ask spreads. The temporal consolidation of trading is consistent with the theoretical model of Admati and Pfleiderer (1988) in which both insiders and a subset of liquidity traders are given discretion over the timing of their trades. Liquidity traders prefer to trade when the market is deep (when trades have little impact on price). This in turn gives liquidity traders strong incentives to trade together. The insider also prefers to trade when the market is deep. The tendency of spreads to narrow over time is also observed in another dealer market experiment (Lamoureux and Schnitzlein (1997)). This type of disequilibrium behavior is roughly consistent with the ad-hoc price adjustment rule posited by Bulow and Klemper (1994).

Most insider trades occur late in the trading period (Figure 3.1): on average, one-quarter of insider volume is not transacted until over two-thirds of the trading interval has expired. Temporal patterns in insider trading almost exactly match patterns in liquidity trading, with the correlation between the time of insider trades and liquidity trades 87% ($p < 0.01$). The temporal concentration of trades is consistent with the theoretical model of Admati and Pfleiderer (1988) in which both insiders and a subset of liquidity traders are given discretion over the timing of their trades. Unlike the model, most trading is concentrated late in the trading period. This important difference is due in part to the complete absence of ex ante constraints on when liquidity traders must trade in the experimental markets. With absolute discretion over the timing of trades liquidity traders prefer to trade late when bid-ask spreads are more closely centered on intrinsic value.

The insider infrequently engages in unprofitable trades (3.7% of all insider trades). The majority of these trades are in the option in the opposite direction of the insider's information. This is evidence that these trades are intended to mislead dealers, since in most trading periods this is the least unprofitable trade when an insider trades in the opposite direction of her information.¹⁹ At the time of most insider trades (72.6%) the insider can trade profitably in either the stock or option market. In most of these cases (86.3%) the insider chooses to transact in the market that at that instant offers the more profitable trading opportunity. When the option is in-the-money, insiders make 51.4% of their trades and earn 44.0% of their trading profits by transacting in the option. When the option is out-of-the-money, insiders make 34.5% of their trades and earn only 8.8% of their profits from trading in the option: the option market's role in the price discovery process is more limited in this case.

¹⁹ Bloomfield and O'Hara (2000) is an experimental study that investigates whether transparent markets can survive in the presence of less transparent markets. In their design, both a liquidity trader and an insider may conceal their trades by trading with a low transparency dealer. They conclude that strategic trading does not play an important role in their markets and traders are unwilling to pay a premium to have their trades concealed, although there is evidence that insiders sometimes trade against their information with transparent dealers in order to fool the market. The strategy space is rich in the sense that traders can choose among multiple trade sizes, however they have limited ability to choose the timing of trades: all trades are executed simultaneously at one of eight trading rounds each period, and therefore time per se does not play an explicit role.

We decompose insider profits by regressing insider profits on a constant, an indicator that takes on the value of one when the liquidation value of the option is positive (ITM), the distance of the stock value from its unconditional expectation (DEV), the number of liquidity trades in the stock (LIQ(S)), the number of liquidity trades in the option (LIQ(O)), the product of the net liquidity shock in the stock and the difference between the stock value and its unconditional expectation (COR(S)), and a similar statistic for the option (COR(O)).

$$\text{PROF}_{\text{INS,STOCK}} = b_0 + b_1\text{ITM} + b_2\text{DEV} + b_3\text{LIQ(S)} + b_4\text{LIQ(O)} + b_5\text{COR(S)} + b_6\text{COR(O)} \quad (2)$$

The results (Table 3.4) indicate that insider profits are increasing in the distance between the stock value and its unconditional expectation ($p < 0.01$). Profits are also higher when the net liquidity shock in the stock is in the opposite direction of the insider's information: the insider can more easily disguise her information while making profitable trades ($p < 0.01$). A similar effect obtains for the net liquidity shock in the option ($p = 0.06$). When the end-of-period option value is positive the insider earns lower profits in the stock ($p < 0.01$). This is due to a greater number of profitable trades in the option (which stock dealers learn from). A regression identical to (2) but where the dependent variable is insider profits in the option (Table 3.4) indicates insider profits in the option are higher when the end-of-period option value is positive ($p < 0.01$).

Aggregate insider profits (profits derived from trading in the stock and the option) are higher when the option is out-of-the-money. Using session level data (Table 3.3) the difference is not significant ($p = 0.32$ and $p = 0.29$ for the parametric and nonparametric test respectively), but the analysis in Table 3.4 reveals a significant difference ($p = 0.02$). If real, this result would be puzzling, since the insider is free to forego option trading. It is driven however by two unusual trading periods when the intrinsic value of the option is zero, and is not a robust result²⁰.

²⁰ In both of these trading periods, there is a large positive net liquidity shock. In one case, eleven consecutive liquidity trades in the stock and the option at the beginning of the period are purchases. Interspersed among these trades are additional (unusual) buy orders initiated by dealers that move price further away from intrinsic value. When the insider finally begins to trade, price moves very slowly toward intrinsic value: the dealers are apparently convinced by the initial 18 consecutive buy orders that the stock value is above its unconditional expectation. In the second case, a single dealer lowers the inside bid unusually slowly in the face of a large

3.2.4 Liquidity trader profits and behavior

Liquidity traders are free to choose the timing of trades, but the exogenous liquidity shocks are enforced by penalties that make liquidity demands (by the end of a market period) perfectly inelastic. Each liquidity trader receives an independent shock in the stock and the option. Over the 62 market periods, the liquidity traders never incurred a penalty for failing to exactly fulfil liquidity requirements.

As noted above, liquidity traders tend to trade late: 86.3% of liquidity trades are in the second half of the trading interval, when spreads are narrower. Requirements to sell options are on average satisfied earlier than requirements to buy options and to buy or sell units of the stock. This may be due to the fact that with the highly skewed unconditional option value distribution, price discovery is on average easiest for option sell orders. The desire of liquidity traders to trade near the unconditional asset value distribution therefore dictates earlier liquidity trading in this case, although two-thirds of these trades are still in the second half of the trading interval.

sell order imbalance. In the first case, insider profits exceed their mean by 3.9 standard deviations, and in the second by 4.8 standard deviations. The high insider profits in these two periods explain the significance of the ITM coefficient estimate, but seem unrelated to the relationship between the level of insider profits and the presence of a positive intrinsic value option. We verify this by repeating the regressions from Table 3.4 but flagging these two periods with an indicator variable that takes on the value of one in these two periods and zero otherwise. The ITM coefficient estimate becomes insignificant in the aggregate insider profits regression. Importantly, other qualitative results in the other eight regressions in the table are unchanged.

Session	Periods	Stock profits			Option profits			Total profits			
		Liq.	Stock		Liq.	Option		Liq.	Stock	Option	
		Insider	Traders	Dealers	Insider	Traders	Dealers	Insider	Traders	Dealers	Dealers
Panel A: Intrinsic value of the option is zero											
1-1	5	41.1	-99.7	58.6	7.5	-10.7	2.7	48.6	-110.4	59.1	2.7
1-2	4	46.5	-38.1	-9.2	4.1	-8.7	4.6	50.5	-46.8	-9.2	5.5
2-1	4	82.9	-100.1	24.5	15.2	-6.9	-5.8	98.1	-107.0	22.0	-13.0
2-2	6	60.7	-86.9	26.4	7.3	-1.1	-1.1	67.9	-88.0	21.3	-1.3
3-1	3	42.6	-77.1	41.8	2.3	-9.2	4.1	44.9	-86.2	44.5	-3.2
3-2	5	43.6	-6.2	-35.7	1.1	-4.5	2.8	44.7	-10.7	-35.1	1.1
4-1	5	73.7	-63.1	-18.8	5.1	-2.3	-1.6	78.7	-65.4	-20.0	6.7
4-2	5	14.6	-78.9	72.6	-1.9	-5.2	7.2	12.7	-84.1	72.6	-1.1
Avg		50.7	-68.8	20.0	5.1	-6.1	1.6	55.8	-74.8	19.4	-0.3
Panel B: Intrinsic value of the option is positive											
1-1	2	53.0	1.9	-57.4	21.0	-6.3	-14.6	74.0	-4.4	-57.6	-12.1
1-2	3	0.3	-33.0	34.7	20.4	-22.5	2.2	20.7	-55.5	34.7	0.2
2-1	4	27.5	-29.8	-6.9	19.2	-15.5	-4.4	46.7	-45.3	-6.1	4.8
2-2	2	17.0	-5.0	-12.0	69.3	-18.8	-50.6	86.3	-23.8	-12.0	-50.6
3-1	5	27.9	-39.7	1.1	9.3	-15.1	22.4	37.2	-54.8	-15.5	33.1
3-2	3	6.7	-26.2	21.8	6.2	-2.4	-4.0	12.9	-28.7	22.1	-6.3
4-1	3	18.0	-160.6	144.5	13.0	-21.9	12.9	31.0	-182.5	140.5	11.0
4-2	3	31.1	-15.2	-15.7	6.0	-21.3	15.3	37.1	-36.5	-15.7	15.0
Avg		22.7	-38.5	13.8	20.6	-15.5	-2.6	43.2	-53.9	11.3	-0.6
Difference		28.0	-30.3	6.3	-15.5	9.4	4.2	12.5	-20.9	8.1	0.3
p-value t-test		3.0%	23.2%	85.0%	5.7%	2.2%	59.9%	30.6%	41.8%	80.7%	97.6%
p-value rand. test		3.9%	21.6%	86.6%	0.0%	3.0%	66.4%	27.7%	40.9%	78.0%	96.5%

Table 3.3. Profitability by trader type and the intrinsic value of the option.

This table reports profitability by trader type for each of the eight sessions. Panel A is based on periods where the intrinsic option value is zero, Panel B on the periods where it is positive.

Market	Trader Type	constant	ITM	DEV	LIQ(S)	LIQ(O)	COR(S)	COR(O)	Adj R2
Stock	Insider	9.91 (0.30)	-36.20 (-3.34)	3.93 (2.72)	-3.39 (-1.50)	8.44 (2.11)	-0.23 (-3.15)	-0.46 (-1.96)	28.50%
Stock	Liq. traders	3.53 (0.19)	25.12 (3.10)	0.90 (2.46)	-10.45 (-4.41)	0.62 (0.13)	0.70 (6.96)	0.06 (0.37)	65.60%
Stock	Dealers	-13.43 (-0.43)	11.07 (0.79)	-4.83 (-3.38)	13.84 (3.84)	-9.06 (-2.00)	-0.47 (-3.42)	0.40 (1.54)	46.90%
Option	Insider	-7.25 (-1.04)	11.55 (3.44)	0.83 (2.68)	0.89 (0.93)	-0.43 (-0.24)	-0.04 (-1.28)	-0.18 (-2.22)	20.10%
Option	Liq. traders	0.87 (0.18)	-11.01 (-5.49)	0.02 (0.12)	0.66 (1.34)	-3.23 (-2.20)	0.00 (-0.15)	0.29 (5.42)	42.10%
Option	Dealers	6.38 (0.90)	-0.54 (-0.16)	-0.85 (-2.35)	-1.55 (-1.71)	3.65 (1.97)	0.04 (0.95)	-0.11 (-1.44)	8.70%
Stock & Option	Insider	2.65 (0.08)	-24.65 (-2.44)	4.76 (3.63)	-2.50 (-1.11)	8.01 (2.06)	-0.27 (-2.87)	-0.63 (-2.94)	32.40%
Stock & Option	Liq. traders	4.40 (0.21)	14.11 (1.53)	0.92 (2.27)	-9.79 (-3.64)	-2.61 (-0.48)	0.70 (6.85)	0.35 (2.28)	59.50%
Stock & Option	Dealers	-7.05 (-0.22)	10.53 (0.73)	-5.68 (-4.19)	12.29 (3.20)	-5.40 (-1.00)	-0.43 (-2.71)	0.29 (1.02)	40.60%

Table 3.4 *Determinants of profitability*

In order to examine the determinants of profitability by trader type we estimate the following model with GMM (t-statistics that are consistent in the presence of heteroskedasticity and autocorrelation are reported in parentheses). Each of the 62 observations summarizes the random draws (asset values and liquidity shocks) for a single market period.

$$\text{PROF}_{\text{type, market}} = b_0 + b_1\text{ITM} + b_2\text{DEV} + b_3\text{LIQ(S)} + b_4\text{LIQ(O)} + b_5\text{COR(S)} + b_6\text{COR(O)}$$

Variable definitions are as follows. $\text{PROF}_{\text{type, market}}$ is the per period profit of all subjects of the same type (insider, liquidity traders or dealers) in a specific market (stock, option or both). ITM is an indicator variable that takes on the value of 1 when the intrinsic value of the option is positive and 0 otherwise. DEV is the absolute difference between the stock value and its unconditional expectation of 100. LIQ(S) is the number of liquidity trades in the stock. LIQ(O) is the number of liquidity trades in the option. COR(S) is the product of the net liquidity shock in the stock and the difference between the stock value and its unconditional expectation (a measure of the extent to which liquidity traders trade in the direction of the insider's information). COR(O) is the product of the net liquidity shock in the option and the difference between the stock value and its unconditional expectation.

Liquidity traders suffer large significant losses in both the stock and option market (Tables 3.1 and 3.3). As noted in Section 3.2.1, average liquidity trader losses in the stock when the option intrinsic value is positive are lower than when the option intrinsic value is zero, although this difference is only marginally significant. In order to control for other determinants of liquidity trader profitability, we regress liquidity trader profits in the stock on the liquidity shock variables, the extremeness of the asset draw and the positive option intrinsic value indicator variable (Table 3.4).

Overall liquidity trader profits are decreasing in the magnitude of the liquidity shocks, but increasing in the correlation between the net liquidity shock and the insider's information: when the liquidity traders are required to trade in the direction of the insider's information, their profits are higher. Both of these effects are highly significant ($p < 0.01$). The extremeness of the asset value draw is also significant ($p = 0.02$). Finally, when the option's intrinsic value is positive, liquidity trader losses are lower ($p < 0.01$). This is because liquidity traders trade at prices closer to intrinsic value in this case.

3.2.5 Dealer profits and behavior

Recall that the three stock dealers are required to make a market in the stock. They are also permitted to transact against quotes submitted by other dealers in the stock and the option. The option dealers participate in 6.7% of the transactions in the market for the stock and the stock dealers participate in 11.0% of the transactions in the option.

Since both stock and option dealers are permitted to trade in either asset we define dealer profits in the stock to be the sum of all stock dealers' profits in a market period, with option dealer profits similarly defined. Weighting each of the sessions equally, stock dealer profits are about half the magnitude of insider profits, and significantly greater than zero ($p = 0.03$ for both the t-test and the randomization test). Option dealer profits are positive, but lower, and are not significantly different from zero. In order to investigate the determinants of dealer profitability, we estimate the regression from (2) and report the results in Table 3.4.

The results are similar when we use either stock market dealer profits or the sum of stock and call option dealer profits (aggregate dealer profits) as the dependent variable. Dealer profits are decreasing in the extremeness of the asset value draw, and increasing in the magnitude of the stock liquidity shock. When stock dealer profits is

the dependent variable the negative coefficient on the COR(S) variable indicates that when the liquidity traders trade in the direction of the insider's information, stock dealer profits are lower. This is intuitive since in this state dealers take fewer highly profitable liquidity trades that are in the opposite direction of intrinsic value. When aggregate dealer profits is the dependent variable, the coefficient estimate for the COR(S) variable decreases in absolute value. This is because when the stock liquidity traders tend to trade in the direction of the insider's information, option dealers learn more rapidly from trades in the stock and earn higher profits. In general, dealer profits are not related to whether the intrinsic value of the option is positive. This is evidence that the dealer competitive dynamic accounts for the informational content of the order flow as a function of trading patterns (that depend on the intrinsic value of the option).

3.2.6 Informational linkages between the stock and option markets

In the previous sections we document the impact of the intrinsic value of the option on price errors, volatility and profitability by trader type. In this section we examine the informational linkages between the markets that give rise to these effects.

As noted earlier, the insider trades aggressively in both the stock and the option. This is consistent with the mixed strategy equilibrium in Easley, O'Hara, and Srinavas (1998), and indicates that a trade in one market will have informational implications for the other market. We document the informational linkages by examining changes in bid-ask spread midpoints induced by transactions (Table 3.5). We define a quote midpoint price change as the difference in the midpoint of the inside bid and ask after a transaction and before a subsequent transaction. If there are multiple quote changes (by one or more dealers) that affect the inside quote, they are summed together.

Average quote changes are consistent with a strong informational linkage between the two markets. Quote midpoints in both markets are revised upwards after a stock or option transaction at the ask, and downwards after a stock and option transaction at the bid.

Transaction induced quote changes are not symmetric across the stock and option. First, a transaction induces larger quote changes in the stock than in the option. This is expected because the option's density covers only half the range of the stock's density. Second, the average responsiveness of stock quotes in absolute value to an

option purchase (L\$1.16) is 68% larger (though not significantly, $p=0.12$) than the responsiveness to an option sale (L\$-0.69). The skewness of the option value distribution provides a logical explanation for this effect. In the beginning of the trading period, the insider can profitably buy options only when the true stock value is sufficiently larger than the unconditional expectation of 100, whereas the insider can always profitably sell options when the true stock value is 100 or lower. An option purchase therefore signals on average a more extreme stock value than an option sale, and this market feature seems to be well understood by the stock dealers.²¹ Third, transactions in the option lead to larger quote changes in the option than in the stock market; a phenomenon consistent with the lower liquidity shocks (and hence greater informational content) of trades in the option.

²¹ On average, opening asks in the stock and call respectively are L\$110.43 and L\$10.42. Average opening bids are L\$91.52 and L\$1.34.

Panel A				
Changes in stock quotes in response to a stock or call transaction prior to the subsequent transaction (stock or option)				
	<u>Stock sell</u>	<u>Stock buy</u>	<u>Option sell</u>	<u>Option buy</u>
# Positive stock changes	45	311	33	72
# Negative stock changes	339	37	86	20
Average	-1.12	0.85	-0.69	1.16
Standard Error	0.14	0.09	0.18	0.24

Panel B				
Changes in option quotes in response to a stock or call transaction prior to the subsequent transaction (stock or option)				
	<u>Stock sell</u>	<u>Stock buy</u>	<u>Option sell</u>	<u>Option buy</u>
# Positive option changes	36	65	24	187
# Negative option changes	119	68	229	33
Average	-0.44	0.20	-0.61	0.63
Standard Error	0.10	0.11	0.08	0.09

Table 3.5 *Linkages between the Stock and Option Markets*

This table documents the informational linkages between the option and stock market by reporting the responsiveness of quotes to trades in both markets. We report changes in the midpoint of “inside” bid-ask spread (the average of highest bid and lowest ask). Panel A summarizes changes in stock quotes that occur after a transaction but before a subsequent transaction. The stock quote changes are broken out as a function of the direction of the transaction (buyer or seller initiated), and the market in which the transaction was made (stock or option). Panel B summarizes changes in option quotes. Both stock and option quotes change in the direction consistent with dealers in one market updating their beliefs with respect to intrinsic value on the basis of transactions in the other.

3.3 Discussion

When the intrinsic value of the option is zero, the insider concentrates her trading activity in the stock. Stock dealers update their beliefs primarily on the basis of stock market order flow and prices gradually converge to the informationally efficient level.

When the intrinsic value of the option is greater than zero, the insider splits her trading activity between the stock and the option. In this case stock dealers learn from both stock market order flow and price discovery in the option. This leads to a significant increase in stock market informational efficiency relative to the case where the intrinsic option value is zero (and most price discovery occurs in the market for the stock). Importantly, differences in opinion among stock and option dealers (and the insider response) are particularly informative, and speed the price discovery process. This effect is not a feature of theoretical models. Although the concentration of insider trading in the stock depends on the ex post moneyness of the option, the insider trades in both assets in both states, as in the mixed strategy equilibrium derived by Easley, O'Hara, and Srinivas (1998).

These results highlight the mechanisms by which the introduction of a traded option can improve the market quality of the underlying asset, and are consistent with the theoretical result of John, Koticha, Narayan, and Subrahmanyam (2000): "...even though the addition of option trading enhances the ability of informed traders to disguise and profit from their trades, the informativeness of the trading process is greater because the market can now infer private information from two sources – order flow in the stock *and* option markets." We believe it is noteworthy that we find support for this result in a richer setting than their model since we allow both insiders and liquidity traders to be strategic. We also show the exact mechanism by which this occurs in a dynamic setting. Unlike the aforementioned model, we do not find evidence in support of bid-ask spreads increasing in the presence of options.

Biais and Hillion (1994) focus on the introduction of a nonredundant option that completes the markets and find ambiguous consequences for the informational efficiency of the market. In their single period model there are three states and three different types of liquidity traders with state-dependent endowments that trade in order to hedge their risk exposures. The introduction of the option can reduce the informational efficiency of the market by enlarging the set of strategies the insider can

follow. This makes it more difficult for the dealers to interpret the informational content of trades.

Although our setting is very different, we do not find evidence for the intuition that a larger strategy space for the insider will reduce informational efficiency. This is most likely due to our dynamic setting with exogenous liquidity shocks in both the stock and the option: the insider uses the cover of liquidity trading to gradually but fully exploit profitable trading opportunities. Since the insider trades more aggressively in the call when it has a positive value, we find that the required number of trades in the stock to reach a given level of efficiency is less than when it has no value. This sharing of price discovery strongly suggests a gain in price efficiency relative to a case without a traded option.

Unlike the theoretical results of Back (1993), our results suggest a decline in the average level of stock volatility as a result of option introduction. A correlated asset in our experiments allows dealers to set their prices with greater precision, because they learn from trades in the other market. We speculate that an important difference between the model and the experimental design that helps account for this result is the presence of strategic liquidity traders in the experiment: in Back's model the liquidity trader arrival process is a Brownian motion and trading is continuous.

3.4 Conclusion

We analyze the informational linkages between a stock market and a traded option by performing a controlled experiment. This allows the observation of all information sets and all actions in a setting based on the Kyle (1985) framework, but beyond the reach of tractable modeling. We examine the hypothesis that the presence of an option improves the market quality of the underlying asset by permitting the effective sharing of price discovery across markets.

We find that an insider trades aggressively in both the option and the stock, with most trades directed to the asset that affords the most profitable trading opportunity. This leads to price discovery occurring in both markets, and hence important feedback effects: trades in the stock market imply quote revisions in the options market and vice versa. We believe this result sheds light on why most empirical studies find an improvement in market quality after the introduction of traded options.

The focus of most related theoretical and empirical literature concerns the effect of the option on the time series properties of the price of the underlying asset. We find a significant relationship, but one that varies dramatically with the ex post intrinsic value of the option: when the option is in-the-money the convergence to informationally efficient pricing is more rapid and the volatility of transaction prices is lower. This is due to option trades making a greater contribution to price discovery. Here the linkage between the two markets is direct in the sense that liquidation values are perfectly correlated. When the option is out-of-the-money, price discovery in the option only helps truncate the stock value distribution. The dependence of market quality in the stock on the option's intrinsic value is very strong, and we thus demonstrate the implications of the presence of a correlated asset for price discovery. Importantly, the fundamental way in which information is extracted from order flow changes in the presence of an option with positive moneyness. The tendency of insiders to trade where the magnitude of the profitable trading opportunity is greatest, provides a richer set of signals to dealers than when there is only a single asset in which the insider can trade profitably. We show therefore that not only does the presence of a correlated asset effectively split price discovery across markets; it also fundamentally changes the process by which conditional expectations are updated. We furthermore show that the less strategic the insider (due to risk-aversion, impatience, or noisy signals), the more powerful we expect this effect to be.

4 Insider strategies with options²²

In this chapter we analyze the same question as in the previous chapter: how does option trading affect market quality in the underlying asset, as measured by speed of convergence to intrinsic value, volatility and spreads. As we outlined in the introduction to both chapters, the amount of empirical work on the impact of option trading is huge. Many of those empirical studies were motivated by the lack of clear guidance that theoretical models provide. With this chapter we aim to produce a better guide: one that clarifies under what conditions a derivative asset improves market quality. For example, in a theoretical paper Biais and Hillion (1994) show that although the option can help avoid a market breakdown, it enlarges the set of trading strategies an insider can follow, and can reduce informational efficiency by making it more difficult for market makers to interpret the information content of trades. The major finding of Back (1993), in an extension of the Kyle (1985) model, is that the listing of an option leaves average volatility unchanged, which contrasts a lot of empirical work.

Instead of experiments we now develop a sequential trade model in the spirit of Glosten and Milgrom (1985) to obtain deeper insights. Of the comparatively small number of theoretical papers, at least two use a sequential trade model in the spirit of Glosten and Milgrom (1985) as well: Easley, O'Hara and Srinivas (1998; further referred to as Easley et al), and John, Koticha, Narayanan and Subrahmanyam (2000; further referred to as John et al). Both models are developed in an asymmetric information setting in which informed and uninformed traders trade in stock and option markets. Easley et al (1998) study whether option markets attract informed trading and whether they incorporate information more quickly than stock markets. They derive that under certain conditions options are attractive to traders with superior information. John et al (2000) focus on the impact of option trading on the market quality of the underlying price process, and the role of margin requirements. Whether options decrease or increase market quality depends on the criterion they employ.

²² This chapter is based on: C. de Jong, 2001, "Informed option trading strategies: the dynamics of the underlying price process", ERIM research paper

We start from the viewpoint that as long as some superiorly informed traders use options, option trades convey information about the underlying. In this chapter we provide two important additions to the existing models, that provide valuable new insights under what conditions and with what mechanisms option trading may improve market quality in the underlying security. First, we extend the single-trade markets to a dynamic multi-trade environment. Second, we analyze market quality under different levels of option leverage, the main distinguishing property of options. Initially, we elaborate on a standard sequential trade model, and show that it is inherently dynamic. Expectations are updated after every trade, which allows us to study a sequence of trades and analyze new and more precise criteria for market quality. Because of the multiple interacting dynamics of the model, we cannot derive the results in closed form, but rely on simulations to report the main dynamics. We show that the focus of the existing sequential trade literature on only first trades leads to the use of inaccurate criteria for market quality. Our analysis indicates that an option may serve as an extra source from which information can be inferred, which speeds up convergence. In trading performance, uninformed traders only benefit from this speedier convergence in well-developed derivatives markets. Uninformed traders are best off in a derivatives market that allows for relatively large (informed) trades, whereas the number of informed traders should be relatively small. This corresponds to well-developed derivatives markets. In terms of price volatility, our model shows that the effect of option trading is rather the opposite: trading in well-developed derivatives markets leads to higher volatility.

The plan of the paper is as follows. In Section 4.1 we describe the model and in Section 4.2 the associated criteria for market quality. In Section 4.3 we analyze the main dynamics of the model and in Section 4.4 we conclude.

4.1 The model

We develop a sequential trade model that is similar in nature to that of Easley et al (1998) and John et al (2000). These two papers add one or two plain-vanilla options to the model of Glosten and Milgrom (1985). In our model trading takes place in a stock and a call option on that stock, and traders choose either of the two assets. Results are qualitatively the same if a put instead of a call option, or both, would be

included. We include only one option type to facilitate the derivations and the insights in the most important mechanisms.

The model is a standard adverse selection model in market microstructure and covers one period. It explains how market mechanisms lead prices to efficient values when some traders have information superior to others. For a discussion of the different assumptions and the resemblance with real world markets, we especially refer to the papers by Glosten and Milgrom (1985) and Easley and O'Hara (1987). The market is quote driven, which means that buyers and sellers trade with a market maker (also referred to as dealer or specialist), who is responsible for providing liquidity through bid and ask quotes. We assume market makers are profit maximizing and risk-neutral and trading takes place for one unit of one asset at a time. Liquidity and inside traders initiate trades. The liquidity traders trade for reasons of liquidity, such as portfolio rebalancing and time-varying consumption and income. We do not further specify their motives for trade, but assume their demand and supply are completely inelastic, so independent of the outstanding quotes, which excludes the possibility of a market breakdown. The informed traders get private and perfect signals regarding the true asset value. They are completely free to engage in trades and will do so in the pursuit of profits. Competition between informed traders causes available profits to vanish quickly and ensures that informed traders maximize profits at every individual trade.

The random variable S represents the intrinsic value of the stock, the random variable C the intrinsic value of the option. The true asset value may be regarded as a value that every market participant agrees upon after all information has become public. The stock value can either take on a low value $X-v$ or a high value $X+v$. These stock values occur with respective probabilities of δ_L and δ_H , whose sum equals unity. The option has an exercise price of X , exactly in between the high and low value, and its value can directly be derived from the value of the stock: $C = \text{Max}[0, S-X]$.

All possible outcomes of a single transaction and their relative probabilities are depicted in Figure 4.1. Here we give an explanation of the trading process. At the beginning of a period, informed traders know whether the stock value is high or low. Next, trading for that period begins. Dealers set quotes to buy or sell during the trading period, execute orders as they arrive, and then revise their quotes. Informed traders optimally buy (stock or option) when the stock value is high, sell when it is low. The probability that they trade the stock is π (π_L when the stock value is low and π_H when it is high), and will be determined endogenously within the model in the

results section. Liquidity traders' behavior on the other hand is completely exogenous. Liquidity traders randomize their trades across stock and option markets, with a propensity for the stock of β and for the option of $1-\beta$, and have an equal probability to buy or sell²³. Informed and uninformed traders anonymously post trades at random intervals in both markets, making it for the dealers a priori impossible to determine whether they trade with a superiorly informed trader or not. The probability that they trade with an informed trader is μ , with an uninformed trader $1-\mu$. To summarize the probability of all different types of trades:

Pr[Insider buys Stock]	$= \frac{1}{2} \mu \pi_H$	(1a)
Pr[Insider sells Stock]	$= \frac{1}{2} \mu \pi_L$	(1b)
Pr[Insider buys Call]	$= \frac{1}{2} \mu (1-\pi_H)$	(1c)
Pr[Insider sells Call]	$= \frac{1}{2} \mu (1-\pi_L)$	(1d)
Pr[Liquidity trader buys Stock]	$= \frac{1}{2} (1-\mu) \beta$	(2a)
Pr[Liquidity trader sells Stock]	$= \frac{1}{2} (1-\mu) \beta$	(2b)
Pr[Liquidity trader buys Call]	$= \frac{1}{2} (1-\mu) (1-\beta)$	(2c)
Pr[Liquidity trader sells Call]	$= \frac{1}{2} (1-\mu) (1-\beta)$	(2d)

²³ Changing the assumption of equal uninformed buying and selling makes the derivation of results more cumbersome, but leaves the results qualitatively unchanged.

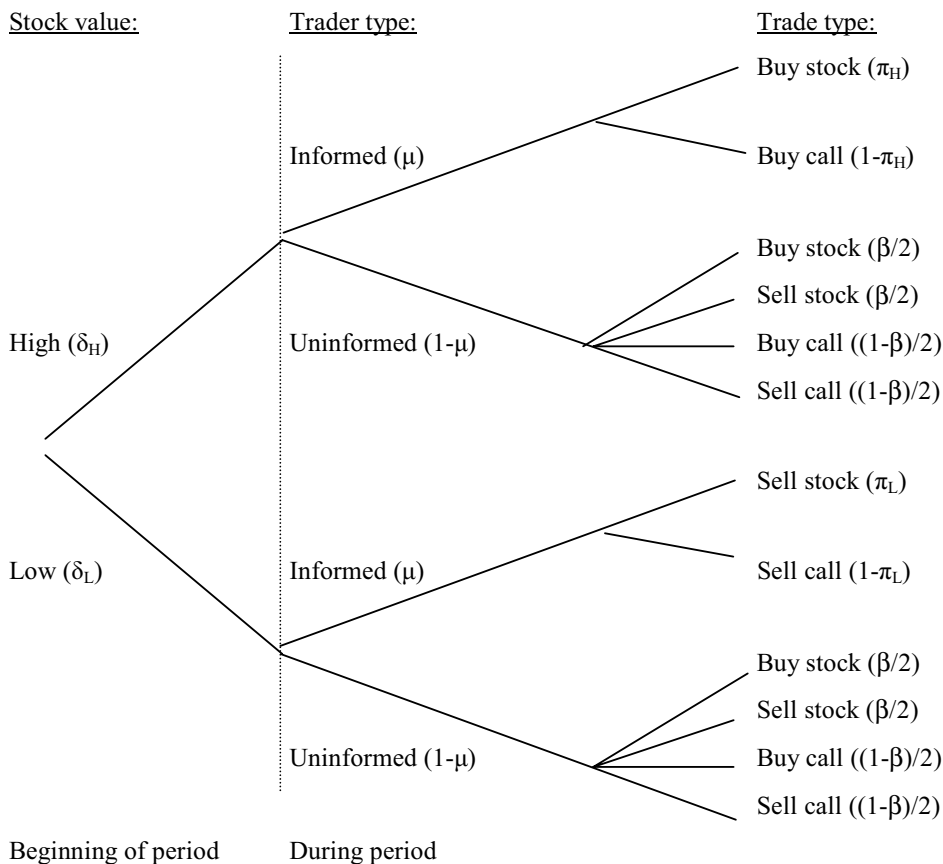


Figure 3.1 *The structure of trading*

The tree diagram shows the structure of trading. Competitive market makers provide liquidity to informed and uninformed traders who may trade in a stock and a call option. The probabilities of the different outcomes of the game are in brackets. The game repeats from the dotted line throughout the period. The game begins at the first node when nature decides whether the true stock value in that period is high or low. Then trading starts and a market maker randomly selects a trader who is allowed to trade. This trader can be informed or uninformed. The variable δ_H is the probability of a high stock value, δ_L the probability of a low stock value; μ is the fraction of informed traders, $1-\mu$ of uninformed traders; π is the fraction of informed traders who trade in the stock; β is the fraction of uninformed traders who trade in the stock; uninformed traders buy and sell with equal probability.

It can be shown that even if an insider is a monopolist, he is better off not entering in unprofitable trades (see for example John et al, 2000). Although unprofitable trades may confuse dealers, the losses incurred cannot be recouped sufficiently to justify such a strategy. If there is at least some probability of uninformed trader activity, the dealers' bids and asks in the stock are strictly in between their minimum and maximum theoretical values. Therefore, when the stock value is high, insiders can profitably buy both assets; when it is low, insiders can profitably sell both assets. As a result, insiders always have two profitable trading opportunities, one in the stock and one in the option.²⁴

From the outset it is unclear what trading strategy would be optimal for the insider. We show that the optimal trading strategy depends on the leverage of the option and explore different levels of leverage. In the first situation the call option gives the right to buy one share of the stock. Although in real world markets an option gives the right to buy several (often 100) shares, this limited leverage might be realistic if we take into account the depth of the market. Since the depth in the option market is generally lower than in the stock market, a representative trade in the option might be for the same number of underlying assets as a representative trade in the stock. A leverage of one may also be realistic from a trading costs point of view. If the order processing costs (excluding bid-ask spread) of one option trade equal that of a trade in the leveraged number of stocks, the actual leverage may also be treated as one. This is the situation explored in John et al (2000). As we will see later in the text, it leads to a relative preference by the insider for the stock market, because the stock has a wider distribution than the option.

In the second situation we allow insiders to trade for a larger (underlying) size in the option than in the stock, as in Easley et al (1998). We pay special attention to the situations where a representative option trade is twice and four times as large as a representative stock trade. This increased leverage is most realistic in modern liquid option markets where order-processing costs are moderate and large option trades do not attract special attention. Obviously, this improves the attractiveness of the option market and leads to an increased number of informed trades in the option.

²⁴ It is informative to list some differences with the experimental design, apart from the important real human behavior a model naturally lacks. The designs are very similar, but in the experiments the asset distribution is more complete, an option never provides leverage (variable in the model), traders are free to enter the market (not randomized, as in several other experiments), and liquidity traders have a fixed (though random) demand for both securities.

As noted earlier, we analyze the model in a dynamic setting, contrary to the existing sequential trade literature. We do so, because a dynamic setting highlights the main effect of derivatives on the underlying asset, which is that dealers in the stock use the information in the derivatives market to set more accurate prices. We introduce dynamics by generating a sequence of trades instead of looking only at opening quotes. An important choice has to be made on where the sequence of trades should stop. The model covers one single period, which could be compared to one trading day in practice, so we should formulate a criterion on how long one period lasts. Important is that the stopping rule is independent of whether the option is traded or not, since the results due to differences in the level of option trading should not be due to differences in the stopping rule. We use the logical assumption that the number of stock trades by uninformed traders is the same for every day, independent of whether options can be traded or not. Implicitly we assume that liquidity based trading in derivative assets is used in addition to trading in the underlying asset, and does not substitute it, as several experimental studies (CBOE, 1975 and 1976) have shown. Under this assumption we can define the end of the day as the moment when uninformed traders have executed a fixed number (N) of trades in the stock. As this does not yet specify the appropriate value for N , the stopping rule is still flexible. We show the results for the situation where $N = 1$, because this is most simple and because higher values for N (we analyzed $N = 2$ and $N = 3$), does not change the paper's conclusions. We thus analyze a sequence of trades in the stock and option that ends with the first liquidity trade in the stock.

4.2 Market quality criteria

Market quality is a rather general term, which may include a number of characteristics of the underlying market. Based on simulated trading sequences of the model described in the previous section we will calculate and analyze different criteria: realized spreads, realized pricing errors and realized price volatility. These criteria are explained later in the text, but since our model is an extension of John et al (2000), we first describe what criteria they use and why they are not appropriate for evaluating efficiency in a multi-asset world.

John et al (2000) study two different criteria, of which the first is the bid-ask spread in the stock at an initial trade. If we define the bid in the stock at the time of the

t 'th transaction (either stock or option) as $B_{s,t}$, the corresponding ask as $A_{s,t}$, then the initial spread $\Delta_{s,1}$ equals:

$$\Delta_{s,1} := A_{s,1} - B_{s,1} \quad \text{Initial stock spread} \quad (3)$$

The authors call their second criterion "the amount of information revealed through trading". It is defined as the ratio of two stock variances: the numerator contains the variance of expectations after the first trade, the denominator the variance of intrinsic values:

$$\eta := \text{Var}(E_1[S]) / \text{Var}(S) \quad \text{Variance ratio} \quad (4)$$

The idea of the variance ratio is that the higher it is, the more information the first trade reveals, the more efficient the market. John et al (2000) find that the inclusion of an option market increases the initial spread (*decreases* market quality), and increases the variance ratio (*increases* market quality). In the discussion of results they place more emphasis on the variance ratio and conclude that an option market improves overall market quality.

We strongly believe that these two criteria are inappropriate for evaluating market quality if trade takes place in more than one asset. The criteria ignore the essential difference an option market introduces, which is that the first trade is not necessarily in the stock, but may be in the option instead. The analysis of opening spreads and opening expectations therefore misses the learning mechanism by which dealers update their quotes. In the results section we will show that the two criteria always yield conflicting outcomes, since a larger initial spread (lower market quality) implies a larger variance ratio (higher market quality). For a sound analysis of market quality, only transaction prices in the relevant asset should be included. That's why we explore more direct methods to evaluate market quality, which are solely based on realized trades in the stock.

We agree with John et al (2000) that an option market may improve market quality, but for a different reason. If an option trade precedes a stock trade, expectations are updated, and bids and asks in the stock are adjusted to the new information. This 'cross-learning' behavior is the main reason that we expect an option market to speed up convergence in the underlying. Ignoring this dynamic effect yields an underestimation of the option's beneficial influence. It should be noted that

John et al (2000) were inspired by Kyle (1985) in the choice of their market quality criteria. In that model however, the difference between expected and realized stock values does not exist, because every trade, and thus every first trade, takes place in the stock. Let us now specify our first two dynamic market quality statistics:

$$\Delta_s \quad := A_s - B_s \quad \text{Realized stock spread} \quad (5)$$

$$PE_s \quad := | P_s - S | \quad \text{Realized stock price error} \quad (6)$$

where P_s is the realized stock price. We call these criteria dynamic, because they require the generation of a sequence of trades. The disadvantage of these dynamics is that the solutions can no longer be derived analytically, but need to be based on numerical simulations. That's why the statistics we report later in the text are the averages of the above statistics for a large number of simulations.

A straightforward way to calculate the above statistics would be at all stock trades, which may occur after a sequence of option trades. An alternative is to analyze the above statistics only at *liquidity* trades in the stock. This can be justified by the notion that new entrants to a market will normally have no specific knowledge about fundamental values. For example, they won't directly bother about the average price errors faced by an insider, at least not beyond the effect it has on their own trades and their own profitability. In order to define the attractiveness of a market, we believe the most logical focus is on uninformed trades. Because the definition of market quality can still be a matter of taste, where it is informative we report market quality both from the viewpoint of an outsider and of all traders (including informed).

Another statistic of interest is the volatility of prices or returns. For a given intrinsic value we calculate volatility as the standard deviation of realized stock prices over a large number of simulations.

$$Vol_s \quad := \sigma(P_s) \quad \text{Standard deviation } (\sigma(.)) \text{ of realized stock price} \quad (7)$$

A large body of empirical work is devoted to the influence of option trading on volatility in the underlying. Volatility may be regarded as an important criterion for market quality, since traders generally prefer few fluctuations in prices. Numerous studies find that option listings cause a decrease in volatility, although in several other

studies the results are mixed or insignificant.²⁵ However, as Skinner (1989) points out, the results should be interpreted with care, because the introduction of an option is not a random event. Market participants and exchanges prefer options on volatile stocks. Reported decreases in volatility after an option introduction may therefore be the result of volatility returning from exceptionally high to normal levels.

4.3 Results

This section reports the effect of option trading on the price errors, volatility and spreads under different levels of option leverage. We analyze the criteria in a state of market equilibrium. For market equilibrium, we first derive the equilibrium dealer quotes for a given insider strategy and then determine what strategy yields an equilibrium outcome. We start in a market where options provide no effective leverage, then continue to higher levels of effective leverage that make options more interesting to trade in.

4.3.1 Equilibrium quotes

Because we assume that dealers are fully competitive and risk-neutral, they set bids and asks in a way that yields zero profits on average. The dealers are uninformed and thus lose on every transaction with a better-informed trader. Uninformed liquidity traders are necessary in this design for the dealers to break even on average. Dealers' quoted bid-ask spread gives them a relative advantage over the liquidity traders, who lose on average.

The zero-profit assumption of dealers can be motivated by the presence of competing dealers or zero entrance costs for new competitive dealers. The dealer sets for example a bid price that equals the stock value conditional on receiving a sales order ($Q = -S$). If he trades with an insider (probability μ), he knows that an insider only sells when the stock value is low ($X-v$, probability δ_L) and when the insider

²⁵ The following find a decrease in volatility: CBOE (1975 and 1976), Trennepohl and Dukes (1979), Skinner (1989), Conrad (1989), Detemple and Jorion (1990), Damodaran and Lim (1991), Kumar, Sarin and Shastri (1998). In the following the results are mixed or insignificant: Klemkosky and Maness (1980), Whiteside, Dukes and Dunne (1983), Fedenia and Grammatikos (1992), Fleming and Ostdiek (1999).

prefers to trade the stock instead of the option (probability π_L). If the dealer trades with a liquidity trader (probability $1-\mu$), he knows this trader sells the stock with probability $\frac{1}{2}\beta$, independent of the true stock value. Using Bayesian inference we obtain:

$$\begin{aligned} B_S &= E[S|Q = -S] \\ &= X - v \cdot \frac{2\mu \cdot \delta_L \cdot \pi_L - (1-\mu) \cdot \beta \cdot (\delta_H - \delta_L)}{2\mu \cdot \delta_L \cdot \pi_L + (1-\mu) \cdot \beta} \end{aligned} \quad (8)$$

Similarly, the dealer sets an ask price that equals the stock value conditional on receiving a purchase order ($Q = +S$):

$$\begin{aligned} A_S &= E[S|Q = +S] \\ &= X + v \cdot \frac{2\mu \cdot \delta_H \cdot \pi_H + (1-\mu) \cdot \beta \cdot (\delta_H - \delta_L)}{2\mu \cdot \delta_H \cdot \pi_H + (1-\mu) \cdot \beta} \end{aligned} \quad (9)$$

The bid and the ask for the option can be derived likewise, keeping in mind that the option value is zero if the stock value is low:

$$\begin{aligned} B_C &= E[C|Q = -C] \\ &= v \cdot \frac{(1-\mu) \cdot (1-\beta) \cdot \delta_H}{2\mu \cdot \delta_L \cdot (1-\pi_L) + (1-\mu) \cdot (1-\beta)} \end{aligned} \quad (10)$$

$$\begin{aligned} A_C &= E[C|Q = +C] \\ &= v \cdot \frac{2\mu \cdot \delta_H \cdot (1-\pi_H) + (1-\mu) \cdot (1-\beta) \cdot \delta_H}{2\mu \cdot \delta_H \cdot (1-\pi_H) + (1-\mu) \cdot (1-\beta)} \end{aligned} \quad (11)$$

4.3.2 Profit maximization without leverage

Throughout the paper we make the assumption that informed traders maximize profits at every individual trade, because they are afraid that otherwise other informed traders will steal away available profits. Furthermore, we assume that informed traders have sufficient investment resources, such that they do not bother about required investments. We believe this is quite realistic, because the information in our model is perfect and because empirical evidence indicates that most informed trading is from large financial institutions with abundant investment resources. In the first part of our analysis we furthermore assume a representative trade in the option is for the same number of shares as a representative trade in the stock. This means that the option market offers no effective leverage. The insider's no-leverage-strategy must logically be of the following form, and depend on the dealer quotes in the market.

Strategy without leverage:

If an insider receives signal L (low stock price), then:

$\pi_L = 0$	if $B_S - X + v < B_C$	always sell the option
$\pi_L = 1$	if $B_S - X + v > B_C$	always sell the stock
$0 < \pi_L < 1$	if $B_S - X + v = B_C$	randomize between stock and option

If an insider receives signal H (high stock price), then:

$\pi_H = 0$	if $A_S - X > A_C$	always buy the option
$\pi_H = 1$	if $A_S - X < A_C$	always buy the stock
$0 < \pi_H < 1$	if $A_S - X = A_C$	randomize between stock and option

We now determine under what conditions each of the above situations hold. Suppose an insider receives the signal L and suppose further that the profit of selling the stock equals that of selling the option. Equating both profits and using the expressions for the bid in stock (8) and option (10), we can derive that insiders transact the stock with probability π_L and the option with probability $1 - \pi_L$, both between zero and one:

$$\pi_L = \frac{\beta}{2\mu \cdot \delta_L} \cdot \frac{4\mu \cdot \delta_L + (1-\mu) \cdot (1-\beta)}{1+\beta} \quad (12)$$

If insiders receive the signal H, and the available profits in both markets are equal, the probability of buying the stock is:

$$\pi_H = \frac{\beta}{2\mu \cdot \delta_H} \cdot \frac{4\mu \cdot \delta_H + (1-\mu) \cdot (1-\beta)}{(1+\beta)} \quad (13)$$

It can be shown that the above π 's exceed β , so insiders have a larger preference for the stock than the liquidity traders have. Please note that if insiders increase their relative preference for the stock (π_H or π_L), the stock spreads widen and the option spreads narrow. This makes it possible to see that if one of the above formulae exceeds one, and so the actual π equals one, the profit of trading the stock is higher than of trading the option. It is also possible to see that the above expressions never equal zero (except for some unrealistic boundary values) and so insiders will never solely trade the option. This is intuitive, because the stock has a larger variability (is more 'information sensitive') and hence cannot offer lower absolute profits than the option.

Static criteria

We start with the static criteria used in other research to show why they yield conflicting results and to obtain first insights. The initial stock spread can be derived analytically, using the expressions for bid and ask (8 and 9), the insiders' equilibrium strategy (12 and 13), and assuming an equal probability of an upward and downward move ($\delta_L = \delta_H = 1/2$).

If $\pi_L < 1$ and $\pi_H < 1$, then:

$$\Delta_{s,1} = (2\mu + (1-\mu) \cdot (1-\beta)) \cdot v \quad (14)$$

The initial spread increases linearly in the probability of an informed trade (μ), and in the distance between the low and high signal (v), and decreases in the

relative preference of the liquidity traders for the stock market (β). If the insiders' preference for trading the stock (π) hits its upper bound of one, the initial spread will be lower than the above expression, because insiders cannot trade the stock as much as they would have wanted. Most interesting is that the initial spread decreases in β , so uninformed traders initially face a higher spread in the stock market the more they trade the option. We obtain a market with only stock trading by setting β equal to its maximum value of one. Then the liquidity traders only trade the stock, and so do the insiders. Using this static criterion we therefore find that the introduction of an option decreases the market quality in the underlying, as does a larger proportion of informed trading.

The second static criterion is the variance ratio. The denominator of the variance ratio, the unconditional stock variance, equals v^2 . The numerator, the variance in expected stock values after one trade, is more complicated. Following a transaction in the stock, the dealers have updated expectations of the stock value equal to the bid (8) or ask (9). Following a trade in the option, the expected stock values are similar to expression (8) and (9), but with π replaced by $1-\pi$, and β replaced by $1-\beta$, the probabilities for the option. If we weigh these updated expectations with the probability of the respective trades, we can derive the analytical expression for the variance ratio (in a mixed strategy)

$$\begin{aligned} \eta &= \mu^2 \cdot \left(\frac{\pi^2}{\mu \cdot \pi + (1-\mu) \cdot \beta} + \frac{(1-\pi)^2}{\mu \cdot (1-\pi) + (1-\mu) \cdot (1-\beta)} \right) \\ &= \mu^2 + \frac{1}{2}(1-\mu)^2 \cdot \beta \cdot (1-\beta) \end{aligned} \quad (15)$$

In contrast to the initial spread, the above criterion reports higher market quality the larger the proportion of informed trading. Furthermore, the above expression is minimal for β equal to zero and one, and has a unique maximum in between (recall that $\pi=\pi_L=\pi_H$ is a function of β and μ). This implies that a market with only stock trading ($\beta=1$) is less efficient than a market with trading in both assets, which is in contrast with the implications from the initial spread.

Dynamic criteria

We derive dynamic criteria, because we believe that static criteria are inappropriate for evaluating the beneficial influence of an option. They ignore the fact that the stock dealers 'learn' from the trades in the option. If the dealers observe for example an option purchase, they know an uninformed trader initiated it with probability $1-\mu$, and that it then does not contain any information about intrinsic values. However, they know that it could also have been an informed trade, and that the stock value must then be high. They update their beliefs according to the following scheme:

$$\text{After a stock purchase: } \delta_{H,+1} = \delta_H \cdot \frac{2\mu \cdot \pi_H + (1-\mu) \cdot \beta}{2\mu \cdot \delta_H \cdot \pi_H + (1-\mu) \cdot \beta} \quad (16)$$

$$\text{After a stock sale: } \delta_{H,+1} = \delta_H \cdot \frac{(1-\mu) \cdot \beta}{2\mu \cdot \delta_L \cdot \pi_L + (1-\mu) \cdot \beta} \quad (17)$$

$$\text{After an option purchase: } \delta_{H,+1} = \delta_H \cdot \frac{2\mu \cdot (1-\pi_H) + (1-\mu) \cdot (1-\beta)}{2\mu \cdot \delta_H \cdot (1-\pi_H) + (1-\mu) \cdot (1-\beta)} \quad (18)$$

$$\text{After an option sale: } \delta_{H,+1} = \delta_H \cdot \frac{(1-\mu) \cdot (1-\beta)}{2\mu \cdot \delta_L \cdot (1-\pi_L) + (1-\mu) \cdot (1-\beta)} \quad (19)$$

After a purchase in stock or option, the probability of a high stock value is revised upwards, and revised downwards after a sale. The larger the insiders' preference for the stock (π_H), and the larger the proportion of informed traders (μ), the more a stock purchase signals a high stock value. The updated beliefs form the basis to generate a sequence of trades. At every point in time, we randomly select a trade and trader type according to the different probabilities. Then we update beliefs and randomly select a new trade and trader type. We continue a sequence of trades till we obtain a (liquidity) trade in the stock, either a purchase or a sale, as we motivated in the 'model'-section. Since there is an infinite number of possible sequences, and

because the beliefs are updated differently in every sequence, we are unable to derive our dynamic market quality statistics theoretically. Therefore, we rely on a large number (one million) of simulations²⁶ to clarify the dynamics.

The parameters X and v are only necessary to scale the stock and option, but do not affect the insiders' strategy or the updating of beliefs. Without loss of generality, we can therefore fix them, for example to 100 and 10 respectively. It is also reasonable that a priori there is no higher probability of an upward move than a downward move, so we keep $\delta_H = \delta_L = 1/2$. The parameters β and μ , that govern the liquidity traders' behavior and the proportion of informed traders, are more delicate, so we will carefully study various values. But first we assume there is an equal proportion of informed and uninformed traders ($\mu = 1/2$), and the uninformed trade as often in the stock as in the option ($\beta = 1/2$).

	Previous trade	δ_H	π	Stock bid	Stock ask	Call bid	Call ask
First trade		0.50	0.83	93.75	106.25	3.75	6.25
Second trade	Stock purchase	0.81	0.77	99.29	108.93	8.13	8.93
	Stock sale	0.19	1.00	91.07	100.71	1.07	1.88
	Call purchase	0.63	0.80	95.36	107.50	5.36	7.50
	Call sale	0.38	0.89	92.50	104.64	2.50	4.64

Table 4.1 *Updating of expectations*

This table reports expectations about the true stock value (δ_H is the probability that the stock value is high), the insider trading strategy (π is the probability that the insider trades the stock), and dealer quotes. Values are reported at the beginning of trading and after one trade has been executed.

In this base case, the initial available profits in stock and option are equal (Table 4.1). That's why insiders start with a mixed strategy, although they prefer the stock five times to the option. Initially, they face a spread in the stock market of 12.50, in the option of 2.50. If a liquidity trader trades before an insider, and trades in the opposite direction of the correct value, an insider trades the stock even more

²⁶ The simulation program (written in Gauss) is available on request.

intensively. However, if he or another trader trades in the correct direction, he subsequently shifts more trading to the option. This happens, because his preference for the stock (π) decreases in the correct expectation. Spreads narrow after the first trade, irrespective of the trade type, because a trade directs the expectations into one direction (the stock distribution becomes skewed). Table 4.1 also shows that a stock trade is followed by a stronger update of beliefs than an option trade. This is due to the relatively higher preference for the stock of insiders than of liquidity traders ($\pi > \beta$).

		$\beta = 0.10$	$\beta = 0.25$	$\beta = 0.50$	$\beta = 0.75$	$\beta = 0.90$	$\beta = 1.00$
Panel A: Market quality statistics based on all stock trades							
$\mu = 0.50$	Spread	2.78	4.04	4.52	4.32	4.05	3.84
	Price error	1.14	2.00	3.08	4.07	4.64	5.03
	Volatility	3.15	3.87	4.34	4.50	4.51	4.50
$\mu = 0.25$	Spread	3.99	4.03	2.80	1.79	1.41	1.21
	Price error	3.46	4.76	6.64	7.94	8.41	8.64
	Volatility	4.58	4.54	4.09	3.35	2.96	2.75
Panel B: Market quality statistics based on the uninformed stock trade							
$\mu = 0.50$	Spread	3.63	5.73	7.20	7.67	7.71	7.67
	Price error	2.27	3.80	5.22	6.17	6.65	6.94
	Volatility	4.93	5.81	5.99	5.76	5.55	5.39
$\mu = 0.25$	Spread	6.68	7.67	7.00	5.81	5.21	4.86
	Price error	5.02	6.55	8.01	8.82	9.12	9.24
	Volatility	5.89	5.58	4.53	3.52	3.08	2.83

Table 4.2 *Market quality with no leverage*

This table reports market quality criteria for various levels of uninformed stock trading. Trades are initiated by an informed trader with probability μ , by an uninformed trader with probability $1 - \mu$. Uninformed traders have an equal probability to buy or sell. They trade with probability β in the stock, otherwise in the option. Insiders maximize profits at every individual trade.

Trading stops after the first liquidity trade in the stock. The statistics in panel A apply to all trades in the stock market, those in panel B to the liquidity trade in the stock market. *Spread* measures the average difference between realized stock bid and ask; *price error* is the average absolute difference between the transaction price and stock value; *volatility* is the standard deviation of stock transaction prices.

Table 4.2 reports the static and dynamic market quality statistics corresponding to this base case and various fractions of uninformed trading in the stock (β). We first focus on the data corresponding to an equal proportion of informed and uninformed traders ($\mu=1/2$). The statistics of panel A are based on the first trade in the stock, those of panel B on the first *liquidity* trade in the stock. All reported values are independent of whether the true stock value is low (90) or high (110). We can infer the following results from Table 4.2 for increasing levels of uninformed trading in the option (decreasing β).

Thanks to the increased coverage of the option market, insiders are able to execute more option trades before the first stock trade. This increased trading activity makes it easier for dealers to form their opinion on the correct stock value, which in turn leads to lower realized price errors in the stock for all trader types.

The effect on stock price volatility is ambiguous: volatility is highest for intermediate levels of option trading. This can be explained by the phenomenon that realized stock prices depend on the number and direction of previous option trades, which vary most for intermediate levels of option trading.

The effect of increased option trading on stock spreads is ambiguous as well. From equation (14) we know that the initial quoted spread increases in the fraction of option trading, but this result does not hold for realized spreads. Realized spreads may increase or decrease with option trading. Both insiders and liquidity traders may face an increased spread. Liquidity traders face only somewhat larger stock spreads when they execute few trades in the option ($\beta=0.9$) compared to when they only trade the stock ($\beta=1.0$). Although the difference is very limited, their modest option trading will harm their performance in the stock market.

To clarify the ambiguous results on realized spread and volatility, we restrict the fraction of informed traders to a more realistic 25% (Table 4.2). It is then easier to see that uninformed traders do not necessarily benefit in the stock market from their activity in the option market. In fact, if uninformed traders form a large proportion of the total population, realized stock spreads (and so stock losses) and volatility increase the more they trade the option. This phenomenon is due to the insiders' strategy, which is aimed at maximizing profits at every individual trade. This strategy leads to a larger preference for the stock than the liquidity traders' preference for the stock. This in turn causes informed traders to reap a relatively large part of all stock trades. Mathematically (see equation 12 and 13), $\pi/(\pi+\beta)$ increases in the fraction of uninformed traders ($1-\mu$) and the amount of uninformed option trading ($1-\beta$).

If the option market provides no effective leverage, the following summarizes the effect of option trading on the price process of the underlying. First, option trading decreases stock price errors. Second, option trading has ambiguous consequences for volatility, realized spreads and losses of uninformed traders. In a market where informed traders form a large part of the population, volatility, spreads and losses may increase or decrease, depending on the exact intensity of uninformed option trading. In a market with a small proportion of informed traders (which is likely to be true in liquid markets), an option market increases volatility, spreads and liquidity traders' losses in the underlying asset.

4.3.3 Profit maximization with leverage

In the above analysis we assumed that the option market provides no effective leverage in the sense that a representative trade in the option is for the same number of shares as a trade in the stock. This situation causes insiders to have a relative preference for the stock, and this preference increases for higher levels of uninformed trading and for higher levels of uninformed trading in the option. Here we relax the leverage assumption and let the parameter γ govern the level of option leverage. Equilibrium bid and ask spreads and optimal insider behavior are now functions of the leverage. They can be obtained by equating the profit in the stock to the profit in the option. Since insiders can trade an option on γ shares, the profit on an option for one share should be equal to $1/\gamma$ the profit on the stock.

A special case is the situation in which the option offers a leverage of two. Then the insiders have the same propensity to trade stock and option as the liquidity traders, as can be shown by equating insider profits in stock and option.

$$\text{Informed stock trading probabilities with leverage } \gamma = 2: \pi_L = \pi_H = \beta \quad (20)$$

With a leverage of two, the insiders in fact mimic the liquidity traders, which makes it hard for the dealers to detect them. This strategy is independent of the quoted bids and asks, the fraction of informed traders and the fraction of uninformed option trading. Because insiders follow the same strategy as the uninformed traders, the probability that a trade in the stock is informed is independent of the uninformed intensity of option trading. Therefore, the initial spread is independent of β and smaller than without leverage:

$$\Delta_{s,1} = \mu \cdot v \tag{21}$$

In Table 4.3 we present the dynamic market quality statistics when insiders mimic the liquidity traders. Although the effect of option trading on volatility is mixed, realized price errors and spreads become smaller the more options are traded, for all trader types. Since insiders and liquidity traders direct an equal proportion of trades to the stock market, the efficiency gain is solely due to the increased learning ability of the dealers, and not affected by differences in preferences between informed and uninformed traders. A leverage of two thus separates the learning effect from the insider strategy effect: an improved market makers' understanding of trades reduces realized price errors and stock spreads.

		$\beta = 0.10$	$\beta = 0.25$	$\beta = 0.50$	$\beta = 0.75$	$\beta = 0.90$	$\beta = 1.00$
Panel A: Market quality statistics based on all stock trades							
$\mu = 0.50$	Spread	1.08	2.03	2.92	3.46	3.70	3.84
	Price error	1.59	2.90	4.03	4.63	4.89	5.03
	Volatility	3.61	4.41	4.63	4.58	4.53	4.50
$\mu = 0.25$	Spread	0.76	0.99	1.12	1.18	1.20	1.21
	Price error	5.56	7.23	8.09	8.44	8.57	8.64
	Volatility	4.87	4.23	3.50	3.06	2.86	2.75
Panel B: Market quality statistics based on the uninformed stock trade							
$\mu = 0.50$	Spread	2.17	4.05	5.85	6.92	7.40	7.67
	Price error	2.13	3.91	5.49	6.36	6.74	6.95
	Volatility	4.39	5.33	5.58	5.51	5.44	5.39
$\mu = 0.25$	Spread	3.02	3.97	4.50	4.73	4.82	4.86
	Price error	5.93	7.70	8.65	9.03	9.17	9.25
	Volatility	5.05	4.37	3.60	3.15	2.95	2.83

Table 4.3 *Market quality when option leverage is two*

This table reports market quality criteria for various levels of option trading. Trades are initiated by an informed trader with probability μ , by an uninformed trader with probability $1-\mu$. Uninformed traders have an equal probability to buy or sell. They trade with probability β in the stock, otherwise in the option. Insiders maximize profits at every individual trade, taking into account that an option trade is for twice as many shares as an option trade ($\gamma = 2$). With this leverage of two they imitate the liquidity traders' preference for the stock relative to the option. Trading stops after the first liquidity trade in the stock. The statistics in panel A apply to all trades in the stock market, those in panel B to the liquidity trade in the stock market. The statistics are based on one million simulations. *Spread* measures the average difference between realized stock bid and ask; *price error* is the average absolute difference between the transaction price and stock value; *volatility* is the standard deviation of stock transaction prices.

		$\beta = 0.10$	$\beta = 0.25$	$\beta = 0.50$	$\beta = 0.75$	$\beta = 0.90$	$\beta = 1.00$
Panel A: Market quality statistics based on all stock trades							
$\mu = 0.50$	Spread	0.03	0.12	0.53	1.58	2.75	3.84
	Price error	1.63	3.02	4.19	4.83	5.03	5.03
	Volatility	3.68	4.55	4.81	4.75	4.63	4.50
$\mu = 0.25$	Spread	0.00	0.00	0.00	0.07	0.54	1.21
	Price error	5.74	7.37	8.03	8.17	8.50	8.64
	Volatility	4.94	4.40	3.98	3.77	3.26	2.75
Panel B: Market quality statistics based on the uninformed stock trade							
$\mu = 0.50$	Spread	0.27	0.82	2.13	4.24	6.11	7.67
	Price error	1.89	3.46	4.86	5.84	6.48	6.95
	Volatility	3.99	4.88	5.24	5.37	5.41	5.39
$\mu = 0.25$	Spread	0.00	0.00	0.02	1.04	3.05	4.86
	Price error	5.73	7.37	8.07	8.35	8.82	9.25
	Volatility	4.94	4.40	3.96	3.75	3.32	2.83

Table 4.4 *Market quality when option leverage is four*

This table reports market quality criteria for various levels of option trading. Insiders maximize profits at every individual trade, taking into account that an option trade is for four times as many shares as an option trade ($\gamma = 4$). See table 4.3 for details of the trading mechanism.

Compared to the situation without leverage, the average spread, losses and volatility in the stock are lower, because the stock dealers do not fear the insiders so much. On the other hand, with this mimicking strategy prices converge more slowly. Dealers have now more difficulty to detect where insiders are trading. Their prices are less accurate and realized price errors larger (but lower than without options)²⁷.

²⁷ One small exception can be detected in Tables 2 and 3 with respect to the smaller price errors. Additional analysis made clear that price errors are only smaller with leverage than without leverage for a combination of unrealistically high proportions of insider trading and uninformed option trading.

We now move on to the general case of leverage. Equating the per trade profit in the stock to the per trade profit in the option and using Bayesian inference we obtain the following probabilities for insiders to trade the stock:

Informed stock trading probabilities with leverage γ :

$$\pi_L = \frac{\beta}{2\mu \cdot \delta_L} \cdot \frac{4\mu \cdot \delta_L + (2-\gamma) \cdot (1-\mu) \cdot (1-\beta)}{\gamma + (2-\gamma) \cdot \beta} \quad (22)$$

$$\pi_H = \frac{\beta}{2\mu \cdot \delta_H} \cdot \frac{4\mu \cdot \delta_H + (2-\gamma) \cdot (1-\mu) \cdot (1-\beta)}{\gamma + (2-\gamma) \cdot \beta} \quad (23)$$

Again, if one of the above expressions exceeds one, insiders only trade the stock. Not surprisingly, insiders trade the stock less intensively when the option provides effective leverage ($\gamma > 1$). It is now also possible for the above expressions to fall below zero, namely if $\gamma > 2$. If option leverage is high enough then insiders may only trade the option. With leverage the traders face an initial spread of:

$$\Delta_{s,1} = (2\mu + (2-\gamma) \cdot (1-\mu) \cdot (1-\beta)) \cdot v \quad (24)$$

The larger the option leverage, the heavier insiders trade the option and the smaller is the difference between initial stock bid and ask.

In Table 4.4 we analyze a market for which the leverage is factor four. Results are similar for other values of γ that exceed two. The results indicate that option trading leads to smaller price errors and smaller realized spreads. Unfortunately, it is hard to predict the impact of option trading on volatility. The effect depends on the exact amount of option trading and the fraction of informed trading.

4.4 Conclusion

The current sequential trade literature with a derivative asset has been concentrated on initial quotes. Under the assumption that option markets provide no effective leverage, opening spreads were found to increase, although trades reveal more information. We use a dynamic setting since the focus on just initial quoted spreads precludes the learning mechanism to take effect. The mechanism by which dealers learn from trades in the other asset to set more accurate prices can only be observed in a sequence of trades. In a multi-trade setting, expectations concerning the true asset value are updated after every trade and price errors, volatility and spreads can be analyzed dynamically by simulation.

Our model clarifies and separates the two mechanisms following the introduction of an option that affect market quality in the underlying stock. On the one hand, stock dealers learn from trades in the option market and set more accurate prices. On the other hand, the proportion of informed trading in the stock is altered depending on the option's effective leverage, possibly reducing market quality. Our dynamic model indicates that option trading reduces price errors in the underlying. The impact is slightly more complicated on price volatility and on the performance of liquidity-motivated traders. The losses by traders with no superior information decrease if the option market provides considerable effective leverage and when there are few informed traders. At the same time, these conditions lead to larger stock price volatility.

Our dynamic, but very stylized model shows that the effect of derivatives markets on the market quality of the underlying asset can only be judged when we know in what market informed traders choose to trade. It is not enough to know that they use derivatives (or not), but also in what proportion to stock market trading. The conditions for a large proportion of informed option trading are best in well-developed markets. In well-developed markets most trading is liquidity-based and informed traders are not easily detected. Moreover, in such markets the liquidity and depth in the derivatives market are relatively large, whereas trading costs are relatively small, which may induce traders with superior information to exploit the leverage of the option market. Our model makes clear that an option market improves informational efficiency in the underlying in terms of reduced price errors. Moreover, in an initial phase a derivatives market may be detrimental to the performance of liquidity traders in the stock market. In the course of its development a derivatives market will

however improve the stock market quality not only in terms of price efficiency, but also in terms of the trading performance of liquidity-motivated traders, though at a price of increased stock price volatility.

5 Conclusion of the first part

The previous two chapters contain the results of two research methodologies to study the impact of an option on the price properties of the underlying asset in a world of asymmetric information. We end this first part of the thesis with a brief review and a comparison of the results.

The differences in methodology have involuntarily created a number of differences in the market designs. Generally speaking, the experiments are a special case of all the different situations we analyzed in the model, but in a more realistic environment with real human subjects, and a more complete asset distribution. For example, in the experiments the options provide no effective leverage, since a trade in the option is for only one unit of the stock. This is a special case in the model, though with a more complete asset distribution and human traders. Similarly, the variations in informed and uninformed stock and option trading that we were able to analyze with the model were impossible in the experiments, due to a lack of time and resources. On the other hand, the experiments provide a more realistic view on reality. For example, they highlight the important and difficult tasks of the dealers, whereas we assume in the model that dealers are perfectly competitive and risk-neutral. Another example is that in the model we assume that each participant knows the exact strategy of all other participants, and behaves accordingly, whereas in the experiments and practice such foresight is unthinkable.

The best comparison between the two methodologies can be made when we consider a situation of no effective leverage in the model. Both theory and experiments indicate that this situation leads to a relative preference of the insider for the stock, because the stock has a wider distribution. Both theory and experiments also predict an important decrease in price errors, because dealers can learn from the prices in the other market to set more accurate quotes.

The methodologies differ however in their implications for some other market quality criteria: stock market volatility, realized spreads and losses incurred by uninformed traders. In the theoretical model the effects on these criteria depend on the fraction of informed trading and the fraction of trades that uninformed traders direct to the option market. Therefore, we have to align these variables with the experimental

design in order to make a realistic comparison. In the experiments we have an insider that accounts for nearly 50% of all trades and uninformed traders transacting one third of their volume in the option. This corresponds to a μ close to 0.50 and a β between 0.50 and 0.75. The experiments find no significant change in spreads, a significant drop in volatility and a significant drop in uninformed stock losses when the option intrinsic value is positive. In a similar situation, the model on the other hand suggests that with option trading realized spreads should increase, but volatility and uninformed losses in the stock should remain more or less constant. We believe that the clear decrease in volatility during the experiments highlights the significant impact that human behavior can have on trading outcomes, and the limitation of a theoretical model. With options having positive intrinsic value, dealers are more confident in their price setting, which has a stabilizing effect on prices. The fact that uninformed losses decrease in the experiments might indicate that informed traders have difficulty in reaping all the benefits options provide. Possibly due to risk aversion, they seem sometimes too eager to trade the option. With such behavior they reveal their intentions too easily and unwillingly support the uninformed traders.

By abstracting from real human behaviour, we believe that the theoretical model somewhat underestimates the benefits of options. In the model all participants draw exactly the same conclusions from each trade. As a result, there are no differences in opinion among stock and option dealers. In the experiments, similar to real-world markets, those differences in opinion (and the insider response) are particularly informative, and speed the price discovery process.

The two studies show the complexity of two interacting markets in a world of asymmetric information. Both the market microstructure as well as human behavior have their share in price formation, but overall lead to lower price errors. Our dynamic, but very stylized model shows furthermore that there may be a trade-off between reduced losses to uninformed traders and increased market volatility. Future research should therefore be devoted to a clarification of this trade-off, if possible in real-world markets, but otherwise in experimental markets to capture the intriguing subtleties of human nature that cannot be ignored in the way markets function.

Regulators that control derivatives markets may use the results of the two studies to better set the standards for derivative markets. Regulators need to weigh the benefits of lower losses for uninformed traders against possibly increased market volatility. They furthermore need to decide on the effective leverage derivative

markets may optimally provide. According to the two studies, the benefits of derivative trading to improved price convergence are however without doubt.

Part II

Empirical studies in derivatives markets

6 Introduction to the second part

Certainly one of the most influential innovations in finance is the famous Black and Scholes (1973) option pricing formula. Each day, many thousands of finance professionals around the world apply this formula and the related Black (1976) formula to price options in their day-to-day business. Its relative simplicity and ease of application are responsible to a large degree for the enormous growth of derivatives trading in the last three decades. In 1997 this was an important motivation for the Nobel Prize Committee to grant the innovators of this formula the Nobel Prize for Economics²⁸.

One of the fundamental assumptions in the Black-Scholes model is that the risk-neutral expected returns on the underlying asset are drawn from a normal distribution. Volatility, measured as the standard deviation of the expected returns, determines the exact shape of this distribution. A good insight in volatility is therefore crucial to calculate proper option values. Since the 1987 crash however, implied volatilities calculated from market option prices have varied over the strike price (or moneyness) and time-to-maturity, instead of being constant as assumed by the Black-Scholes model. The variation indicates that risk-neutral expected returns are not normally distributed.

In this part of the thesis we empirically explore three different option pricing methodologies that deviate from this standard assumption of normally distributed returns. The first two chapters do not aim to price options, but revert the process and infer information from market option prices. The last chapter is a more standard pricing study, but applied to a very non-standard and risky commodity, electricity.

The first chapter presents a methodology whose purpose it is to derive the *risk-neutral* distribution from option prices in a flexible and accurate manner. The exact shape of the implied risk-neutral distribution gives important information that can be used for pricing other options on the same underlying asset, for comparing

²⁸ Nobel Prize winners were professors Scholes and Merton; their colleague Black died some years before.

options on different assets and for closely monitoring changes in the markets perception of the underlying price process.

In the second chapter we explore a method to infer from option prices a forecast of the actual price process. Instead of only a cross-section of option prices, the method employs options with different maturities. It yields a price process and not just a distribution at a single point in time. Combined with a risk premium parameter obtained from actual prices of the underlying security, we convert the risk-neutral process in an actual price process and assess its ability to forecast short-term volatility. Compared to the approach in the previous chapter, we lose some of the flexibility by fitting one process through options with different maturities. The extension from the risk-neutral world to the actual world enlarges however its scope of application.

In the third chapter we shift attention from stock indices to electricity, a recently liberalized and deregulated commodity in many countries around the world. Most of the transactions in this market are still for the physical asset, but futures and other financial derivatives steadily increase in trading volume. One of the best examples of this is the trading volume on the Scandinavian power exchange Nord Pool (see Figure 6.1), which nearly doubles year on year. Moreover, a large proportion of electricity end user contracts contain embedded options, such as caps and swings, and several power generators may be valued with real option theories. Since electricity cannot be stored efficiently, and prices violate normality assumptions severely, pricing those electricity derivatives poses however a formidable task. In the third chapter we provide a framework based on regime switches to price a category of electricity derivatives, namely options on spot prices, and apply it to the Dutch power market.




Figure 6.1 *Trading volume Nord Pool financial market*

This figure shows the total annual volume and its contractual value on the Scandinavian power exchange Nord Pool. The value of 1 Norwegian Kronor (NOK) equals approximately € 0.135 (October 1, 2002). Source: Nord Pool annual report 2001.

7 The skewed-t implied distribution model ²⁹

In the past two decades the Black-Scholes (1973) option pricing model has become the widely accepted standard to value a wide range of derivative securities. Despite its ease of application however, the model exhibits some well-known deficiencies, including its assumption of continuous and costless hedging opportunities and its assumption of security prices following a Brownian motion. These deficiencies become most clear in the model's inability to price options consistently across strike and maturity. For example, the model frequently misprices deep in-the-money and far out-of-the-money options, as was already documented by Black (1975). In fact, since 1987, implied volatility has been a convex function of strike price and referred to as a 'skew' or 'smile', depending on its exact shape (Bates (1991), Rubinstein (1994), Derman (1999)).

Various changes to the Black-Scholes model have been investigated to account for its biases. They may for example be related to non-hedgable risks in relatively illiquid options. The most popular adaptation is however the use of a different underlying risk-neutral price process or distribution. The choice of this adaptation is supported by empirical distribution analysis that rejects Brownian motion for most traded assets. The skews and smiles are thus explained by the non-normal characteristics of implied risk-neutral return distributions.

Cross sections of option prices have long been used to derive implied risk-neutral distributions. These distributions represent a forward-looking measure of future risk-neutral realizations of the underlying security. Option-implied distributions have the distinct advantage of being based on data from a single point in time, rather than from a historical time-series. As a result, these implied distributions are theoretically more responsive to changes in market's perceptions than are forecasts from historical time series data.

As a drawback, the distribution inferred from option prices is risk-neutral. If the representative investor who determines option prices is not risk neutral and cannot

²⁹ This chapter is partly based on: C. de Jong and R. Huisman, 2000, "From skews to a skewed-t: modeling option-implied returns with a skewed Student-t", ERIM research paper

hedge all exposures (as is typically the case), these distributions may not correspond to the market's actual forecast of the future distribution of the underlying asset. This is a drawback easily neglected by many researchers and may lead to incorrect interpretations about market expectations³⁰. Without specifying risk premia (as we do for example in the next chapter), implied distributions should not be interpreted as unbiased predictors of future distributions.

However, even without the specification of risk premia, the risk-neutral implied distribution contains useful information. First, the exact shape of the implied distribution can be used for pricing options on the same underlying asset, such as options with illiquid strikes and maturity, or otherwise exotic features. Second, the implied distribution can be used for comparing options on different assets. This is especially worthwhile when the underlying assets are regarded as having similar representative investors and thus risk-premia are deemed to be equal. Differences in implied distributions that are not present in actual distributions may hint at anomalous market prices. Finally, implied distributions are an excellent means for closely monitoring changes in the markets perception of the underlying price process³¹. This latter application is especially popular among central bankers for gauging the market's expectations regarding interest rates and exchange rates (e.g. Federal Reserve Bank of Atlanta (1995), Deutsche Bundesbank (1995), Campa and Chang (1995), Leahy and Thomas (1996), Malz (1996 and 1997), Campa, Chang and Reider (1997)). Central bankers and other monetary policy makers use this information among others in assessing monetary credibility, and the timing and effectiveness of monetary operations.

Breeden and Litzenberger (1978) were the first to show how the implied risk-neutral distribution function could be derived from option prices: the densities are equal to the second order derivatives of call option prices with respect to the strike price. Shimko (1993) offers a practical application of this general idea. He proposes to model the volatility smile as a quadratic function of moneyness (to obtain a continuous volatility smile), and then to calculate the second order derivative numerically. This approach is simple and fast, but inaccurate outside the range of traded strike prices, where volatility often becomes unbounded. Other methods

³⁰ Anagnou, Bedendo, Hodges and Tompkins (2001) provide an excellent review.

³¹ The word 'changes' should be stressed, since implied risk-neutral distributions should not be interpreted as absolute forecasts, as we discussed.

construct implied binomial (Rubinstein (1994)) or trinomial trees (Derman, Kani and Chriss (1996), Nagot and Trommsdorff (1999)), or estimate the end-of-term distribution non-parametrically (Aït-Sahalia and Lo (1995), Jackwerth and Rubinstein (1996)). Other methods that are suitable to infer skews and smiles can be based on any option-pricing model that relies on non-normal returns. Examples are the jump models of Merton (1975) and Amin (1993), and GARCH option pricing models (see next chapter).

In this chapter we present a different methodology to infer the implied risk-neutral distribution function from European-style options. We use a skewed version of the Student-t distribution, which is known to provide a good fit to historical returns on many financial assets (see Reiss and Thomas (1997) for an overview). The skew or smile pattern of implied volatility as a function of strike is a direct indication of skewness and excess kurtosis of the implied risk-neutral return distribution. A smile implies fat tails; a skew implies both fat tails and skewness. The skewed Student-t distribution we use, is able to capture these distributional moments, and was first proposed in Fernandez and Steel (1998) and later applied to financial time-series by Lambert and Laurent (2001). The method to obtain skewness is simple; it assigns unequal weights to the distribution on the left and right side of the mode. The advantage of this skewed-t method is that the whole distribution depends on only three parameters, of which two directly control for the levels of skewness and kurtosis. Moreover, the skewed Student-t nests the normal distribution. We can thus easily vary parameters to compare different distributions and use the parameters as inputs to price other options. Other methods to obtain a skewed Student-t distribution are given in Hansen (1994), Theodossiou (1998) and Mittnik and Paoletta (2000).

A method that models skewness and kurtosis even more directly than the skewed-t is the one in Corrado and Su (1996). This model adapts a Gram-Charlier series expansion³² of the standard normal density function to yield an option price formula that is the sum of a Black-Scholes option price plus adjustment terms for non-normal skewness and kurtosis. Although conceptually similar to our approach in the sense that it has separate parameters for skewness and kurtosis, it has one important limitation. Implied distributions often have such pronounced tail-fatness, that the fourth moment is non-existent, as we will show empirically. Estimating implied kurtosis then yields spurious results, whereas our skewed-t method is still able to

³² See for example Jarrow and Rudd (1982)

estimate a parameter (the degrees of freedom) that captures tail-fatness and thus describe the underlying distribution.

In the next section we explain our method, as well as two curve-fitting and an implied trinomial tree method. Section 7.2 starts with an example and provides an empirical analysis of the methods with closing prices of European-style options on the FTSE 100 index from January 1995 to December 1999. We describe how we estimate the models weekly and compare the results. Section 7.3 concludes.

7.1 Methodology

The skewed-t method assumes that the expected risk-neutral returns implied in option prices follow a skewed Student-t distribution. In this section we introduce this distribution and describe how it can be applied to options pricing. Then we describe the other methodologies with which we compare the skewed-t approach.

7.1.1 Skewed Student-t

Consider the probability density function $f(x | \alpha)$ of the central Student-t distribution with α degrees of freedom. It reads:

$$f(x | \alpha) = c(\alpha) \cdot \left(1 + \frac{x^2}{\alpha} \right)^{-\frac{\alpha+1}{2}} \quad (1)$$

where $c(\alpha)$ is a constant that exclusively depends on α and ensures the integral sums to unity. The central Student-t distribution is symmetric with mean equal to zero. For values of α larger than two, the variance is defined and equals $\alpha/(\alpha-2)$. The parameter α is called the number of degrees of freedom and controls the level of tail fatness: the smaller the degrees of freedom, the fatter the tails. The Student-t distribution nests the normal distribution: if α approaches infinity the Student-t converges to the normal distribution.

Fernandez and Steel (1998) describe a very general method to introduce skewness into a symmetric distribution (around 0) by transforming the probability density function by the parameter γ as follows:

$$g(x | \gamma) = \begin{cases} f(x \cdot \gamma) & x < 0 \\ f(x/\gamma) & x \geq 0 \end{cases} \quad (2)$$

The basic idea underlying (2) is simply the introduction of inverse scale factors both sides of the mode. This inverse scaling of the probability density function leaves the unique mode at 0, but enables control of the skewness. The density function is symmetric for $\gamma = 1$, negatively skewed for $\gamma < 1$ and positively skewed for $\gamma > 1$. Furthermore, it can easily be seen that inverting γ produces the mirror image around 0.

$$g(x | \alpha, \gamma) = g(-x | \alpha, 1/\gamma) \quad (3)$$

One of the appealing properties of this inverse scaling method is that the moments can be calculated directly from the moments of the symmetric distribution, and that these moments exist if and only if the corresponding moment of the parent distribution exists. In particular (Fernandez and Steel (1998)):

$$E[x^\delta | \gamma] = M_\delta \frac{\gamma^{\delta+1} + (-1)^\delta \gamma^{-\delta-1}}{\gamma + \gamma^{-1}}, \quad \delta = 1, 2, \dots \quad (4)$$

$$M_\delta = 2 \cdot \int_0^\infty s^\delta \cdot f(s) \cdot ds \quad (5)$$

For the Student-t the latter expression is available in closed form, and some extra algebraic equations lead to relatively simple equations for the mean m and standard deviation s (Lambert and Laurent, 2001)

$$M_\delta(\alpha) = \sqrt{\frac{(\alpha-2)^\delta}{2}} \cdot \Gamma\left(\frac{\delta+1}{2}\right) \cdot \Gamma^{-1}\left(\frac{\alpha}{2}\right) \cdot \Gamma\left(\frac{\alpha-\delta}{2}\right), \quad (6a)$$

$$m = \sqrt{\frac{\alpha-2}{\pi}} \frac{\Gamma\left(\frac{\alpha-1}{2}\right) \gamma - 1}{\Gamma\left(\frac{\alpha}{2}\right) \gamma} \quad (6b)$$

$$s^2 = (\gamma^2 + \gamma^{-2} - 1) - m^2 \quad (6c)$$

where $\Gamma(\cdot)$ is the Gamma function, and $\alpha > \delta$. Note that γ not only influences the skewness, but also the mean and standard deviation of the distribution. In the option calculations we will need to standardize the random draws from the skewed-t, and expressions (6a-c) will then turn out to be extremely helpful.

We assume that the natural logarithm of the risk-neutral return of the underlying asset has a standard deviation of σ and follows the above skewed-t distribution, where γ determines the level of skewness and α the level of kurtosis. The mean of this risk-neutral return should be equal to the risk-free rate (r) minus the continuous dividend yield (d). Under these assumptions the time T price of the underlying (S_T) equals:

$$S_T = S_0 \cdot \exp\left\{\left(r - d - \frac{1}{2}\sigma^2\right) \cdot T + \sigma \cdot \sqrt{T} \cdot \left(\frac{x - m}{s}\right)\right\} \quad (7)$$

where x is a random draw from the skewed-t distribution, whose mean and standard deviation equal m and s . Once the three parameters (σ , γ , α) of the underlying risk-neutral process are known, calculation of European-style option prices is straightforward. For calls (puts), the value of the option is simply the value of the portion of the distribution above (below) the strike price, discounted back to the present by the risk-free rate. In particular, the price of a call option expiring at time T at a strike price of K equals:

$$C = e^{-rT} \cdot E[\max(S_T - K, 0)] \quad (8)$$

where the expectation refers to the risk-neutral distribution. The normal distribution is nested in the skewed-t, since the latter equals the standard normal if $\gamma = 1$ and α approaches infinity. Therefore, the skewed-t option pricing methodology generalizes the standard Black and Scholes (1973) methodology by two extra parameters that

directly account for two more distributional moments³³. Unlike the Black-Scholes formula, there is no closed form for the option price under the skewed-t distribution. As a consequence, the expected value in (8) has to be derived numerically.

Although later in this section we show how to infer the parameters from European-style options, it is worthwhile mentioning roughly how they can be inferred from American-style options as well. For American-style options the relationship between the distribution and the option price is less direct, due to the early exercise premium. The option's value will depend on the entire stochastic process of the underlying, not just the distribution at the option's expiration. This early exercise premium however is subject to a rather strict minimum and maximum (see e.g. Melick and Thomas (1997)). For call options expiring at time T the bounds are:

$$C^u = E[\max(S_T - K, 0)] \quad (9)$$

$$C^l = \max\{e^{-rT} \cdot E[\max(S_T - K, 0)], E[S_T] - K\} \quad (10)$$

The upper bound C^u is the undiscounted European option value, whereas at a minimum the lower bound C^l equals the (discounted) European value. At a maximum the ratio of the upper to the lower bound is thus e^{-rT} . Although these bounds can be remarkably close together for reasonable interest rates, a point estimate is required in the inference of the distributional parameters from American options. To generate such a point estimate, one could weigh the upper and lower bounds by an extra parameter that determines where exactly between the bounds the American option is priced³⁴. So, with one extra parameter any methodology could be applied to American options as well.

7.1.2 Other methodologies

With its separate parameters for the first four distributional moments, we expect the skewed Student-t distribution to be flexible enough to price options across different strike prices. The only way to assess its flexibility however is a careful

³³ As can be seen in Equation (6c) the parameter γ not only influences skewness, but the variance and kurtosis as well. Both are minimal for $\gamma = 1$.

³⁴ Melick and Thomas (1997) use different weights for options in- and out-of-the-money.

comparison to a set of other methodologies. Apart from a comparison with the lognormal distribution (constant volatility), we choose three main reference methodologies with approximately the same degrees of freedom (parameters) as our method: two implied curve-fitting methods plus a trinomial tree. Curve-fitting methods were introduced by Shimko (1993) and since become very popular among academics and practitioners alike. Therefore, it is a natural reference method. Of the various tree approaches that have been proposed for option pricing, a trinomial tree has proven to provide the necessary flexibility to model empirical returns (Bliss and Ronn (1989)) as well as option-implied returns (Derman, Kani and Chriss (1996) and Nagot and Trommsdorff (1999)), and is a natural candidate as well. Below we describe these reference methodologies in more detail.

Shimko (1993) was the first not to model the underlying process or distribution, but the implied volatility curve instead. He fits a curve through the implied volatility curve, translates the implied volatilities into option prices and derives the risk-neutral distribution from the second order derivatives of the call option prices with respect to strike prices, following Breeden and Litzenberger (1978).

$$f_{s_T}(K) = \exp(r \cdot T) \cdot \frac{\partial^2 Call(K, T)}{\partial K^2} \quad (11)$$

Shimko (1993) imposes a quadratic polynomial structure on the implied volatility curve. In the original paper the exercise price serves as independent variable in the quadratic form. This method requires no optimization, and is therefore fast and simple. However, it produces option prices that are inaccurate outside the range of traded strike prices, because implied volatilities either go to plus or minus infinity. Hence, tail behavior of the distribution is hard to evaluate. As a solution, it has been proposed to apply a lognormal structure on the tails by flattening the volatility curves at the endpoints. A less ad hoc solution is that of Malz (1997), who proposes to replace the strike price by the option's delta as independent variable, a statistic that lies in the closed interval from zero to one. In our applications we will use both modifications: the deltas as independent variable, and the (rescaled) strike price³⁵.

³⁵ This rescaled strike price is the moneyness: $Ke^{-rT}/S_0 - 1$

A very flexible approach to price options is by constructing trees. Especially for the valuation of early exercise premia in American-style options a tree approach is often applied, since a well-defined tree allows for efficient backward valuation procedures. In a binomial tree the underlying risk-neutral process may go up or down in each time-step. The most well known binomial tree is probably that of Cox, Ross and Rubinstein (1979), in which the magnitude of a downward move is the inverse of the magnitude of an upward move. This relatively simple process has however proven to provide only limited flexibility (Bliss and Ronn, 1989). That's why the tree process may be extended with a state of no or limited change in the underlying process. This extension enlarges its flexibility considerably and is our motivation to choose an implied trinomial tree as a benchmark for our skewed Student-t approach. In our comparison we apply the same tree as Nagot and Trommsdorff (1999). The trinomial tree approach captures those features of implied distributions that are most prominent in real-world option markets: negative skewness and excess kurtosis. In contrast to some other tree approaches (e.g. Derman, Kani and Chriss (1996) and Jackwerth (1997)) theirs requires less information, while maintaining enough flexibility: the underlying asset can either go up by a factor $u (>1)$, not move at all, or go down by a factor $d (<1)$, with respective probabilities of p_1, p_2 and $1 - p_1 - p_2$. Because the downward move is not restricted to the inverse of the upward move, this tree is non-recombining, which explains part of its flexibility³⁶. For reasonable accuracy the number of time-steps should be at least 15 to 20. Going forward through the tree, the probabilities and terminal asset values can be expressed in terms of the four parameters u, d, p_1 , and p_2 .

7.2 Empirical Results

We start this section with an example to gain some insights in the different methods. We proceed with a description of the FTSE-100 options data for the more rigorous empirical analysis. With the five years of options data we estimate the implied distributions and perform a comparative in-sample analysis to assess the appropriateness of the skewed Student-t formulation.

³⁶ Only one degree of freedom is left unused: the intermediate move i may be free to deviate from 0.

7.2.1 Example

In order to gain insight in our Student-t based model, we start with an example that is presented both by Shimko (1993) and Nagot and Trommsdorff (1999). We compare our results with theirs, as well as the classical lognormal distribution.

Strike price	325	345	360	365	375	385
Option price	66.500	46.000	33.000	27.750	20.125	13.500
Strike price	390	395	400	405	410	425
Option price	9.625	7.250	5.375	3.375	1.875	0.250

Table 7.1 *Shimko example*

This table presents the closing prices of European-style call options on the S&P 500 index with different strike prices on October 21, 1991. The options expire in December 1991 and prices are in US\$. The data are taken from Shimko (1993).

Consider the following European-style call options on the S&P 500 index. The prices are from October 21, 1991. The index value is 390.02, the interest rate 5.03%, the continuous dividend yield 3.14%, and the time-to-maturity 0.16 years (40 trading days). The prices of the call options are listed in Table 7.1. The optimization of the skewed-t method leads to the following estimates for the standard deviation (σ) 15.60%, skewness (γ) 0.523, and degrees of freedom (α) 15.64. The skewness parameter indicates pronounced negative skewness and excess kurtosis, due to a moderate number of degrees of freedom and γ deviating from 1.

	RMSE	Skewness	Kurtosis
Normal	1.506	0.000	3.000
Shimko (K)	0.429	-0.348	3.327
Shimko (delta)	0.819	-0.981	4.122
ITT	0.557	-0.633	3.717
Skewed-t	0.443	-0.227	3.263

Table 7.2 *Results Shimko example*

This table presents the implied distribution characteristics of the logreturns from the example presented in Table 7.1 for five methods. Parameter estimates are obtained by minimizing the root mean squared error (RMSE) between actual and model call option prices, except for Shimko's method. 'Normal' refers to the normal implied distribution, 'Shimko (K)' to Shimko's method regressed on rescaled strike prices and flattening beyond traded strike prices, 'Shimko (delta)' to Shimko's method regressed on option deltas, 'ITT' to Nagot and Trommsdorffs (1999) trinomial tree with 30 time-steps, and 'Skewed-t' to our method.

Based on these data we estimate parameters for five different methods by minimizing the root mean squared error (RMSE). The methods under consideration are the standard normal, the Shimko (regressed alternatively on rescaled strike price and option delta), the implied trinomial tree with 30 time steps (ITT), and our skewed-t. In Table 7.2 we compare the fit as well as the return moments (statistics are based on the logreturns). The fit of the skewed-t method and ITT are of comparable magnitude, as indicated by the root mean squared error between the actual and fitted prices. In this example they clearly improve upon the normal method and (to a lesser extent) the Shimko (delta) method. The main differences appear in the estimates for skewness and kurtosis. By construction, the normal distribution contains neither skewness nor excess kurtosis. Due to the flattening of the tails we employed, the estimates of skewness and kurtosis in Shimko's approach are rather unreliable. The reported statistics of all methods differ quite a bit, although the shapes of the density functions look rather similar, except for the normal (Figure 7.1). The differences in moment statistics are likely to be due to the behavior in the tails, which is hard to judge by eye. What this example therefore clarifies, is that the option prices do not completely pin down the underlying distribution, a fact that Melick and Thomas

(1997) provides an excellent graphical example of. This motivates our choice for a parsimonious distribution that has proven to provide a good fit to return-data.

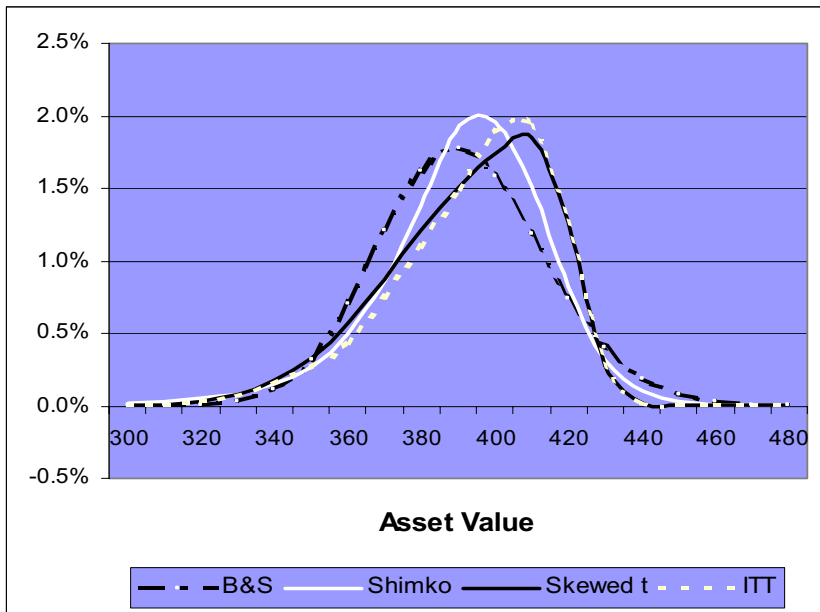


Figure 7.1 *Implied distributions, Skimko example*

This graph shows the implied probability densities, obtained by minimizing the root mean squared error between actual and model option prices, as presented in Table 7.1. 'B&S' refers to the normal (Black-Scholes) method, 'Shimko' to the quadratic curve fitting method with rescaled strike prices, 'Skewed-t' to our skewed Student-t method, and 'ITT' to the trinomial tree approach.

7.2.2 Description of the data

Daily closing prices on the FTSE-100 (Financial Times Stock Exchange 100 Index) index options, traded on Liffe (London International Financial Futures and Options Exchange) are used in an attempt to assess the different methods more rigorously than the above example. We use data from January 1995 till December 1999. The underlying value of the options is the future on the FTSE-100 index with the same expiry date. The FTSE-100 index consists of the largest 100 UK companies by market value, and is the leading indicator for stocks that trade on the London Stock Exchange.

The FTSE-100 options are European-style index options that expire on the 3rd Friday of the expiry month. Over the past ten years, volume has continuously been growing to a daily level of nearly 25.000 contracts, or £30 billion, at the end of 1999. Trading volume is somewhat higher in calls than in puts. Most active trading is in (close to) at-the-money series and short-term maturities. For example, nearly half the volume in the 1995-1999 period is concentrated in the series that expire within one month (20 trading days), whereas only 12% is concentrated in maturities of over three months (60 trading days). Similarly, 62% of the trades are at strike prices within 5% from the current futures price, whereas only 15% is at strike prices outside 10% of the futures prices. There is an apparent difference in strike preference between call and put trades. Most call trades are at strike prices close to or above the current futures price, whereas most put trades are on the other side of the strike spectrum. This shows that traders prefer at-the-money and out-of-the-money options: those options that are cheapest and have the strongest option-like characteristics. Measured by trading volume and open interest, the FTSE-100 options and futures markets are currently the most liquid derivatives markets in Europe. Therefore, they are probably the best markets in Europe for testing an option-pricing model.

The data we use in the analysis are from Liffe's CD-Rom "Equity Products End of Day Data", which contains options and futures data on the FTSE-100 and its constituents from March 1992 till December 1999. We use the closing prices on the FTSE-100 index from January 1995 till December 1999, because markets became more liquid over the years. In order to reduce the amount of calculations, we restrict attention to the Wednesdays³⁷. For liquidity reasons, we restrict attention to the

³⁷ On Wednesday 19 May 1999 there was no trade. Hence, we used the day afterwards.

options whose strike prices are within a 10% range from the current futures price and with a time-to-maturity of at least 5 and at most 127 trading days (half a year) to maturity. Furthermore, since the option prices have tick size of £0.50, we delete from the sample all quotes smaller than £1.50 to avoid problems due to stale prices.

In the option calculations a risk-free interest rate is required to discount back the possible payoffs. Of course, we could rely on some real-life interest series, but then we would need to interpolate or extrapolate interest rates from different maturities to get interest rates with maturities that match those of the options. Shimko (1993) presents an elegant method to circumvent this problem. For every maturity series, he estimates the discount factor that produces a put-call relation that comes closest to put-call parity.

$$Call - Put = F + \exp(-r \cdot T) \cdot K \quad (12)$$

We estimate this relation (with F and $\exp(-rT)$ as parameters) by ordinary least squares and thus obtain an ‘implied’ futures price and interest rate. The implied interest rate may be interpreted as the lowest borrowing and lending costs that market participants face. As a final step in our selection process, we now discard the put prices, because the deviations from put-call parity are negligible. This leaves us with 15466 call option prices, on 258 trading days, with at least 6 different exercise prices per maturity, and average time to maturity of 48 trading days.

7.2.3 Estimation results

We start estimation with the skewed Student-t method and then move forward to the other methods. The skewed-t attains a relatively good fit to the option price data (Table 7.5). Average root mean squared error between actual and fitted prices is £0.77, which should be judged in the light of option values averaging around approximately £220 and 16.5 options per maturity bucket on average. Parameter estimates are quite stable over the years, with the exception of volatility. As expected and in line with other research, volatility was considerably higher in the years 1997-1999 than in the two previous years (Table 7.3). The number of degrees of freedom varies widely and assumes sometimes very low levels (below 2) and sometimes very high levels (exceeding 1000). Because the estimates for the degrees of freedom are so much right-skewed we estimated (and report) its natural logarithm instead of the degrees of

freedom itself. The average of this estimate is 2.63, which corresponds to 13.8 as degrees of freedom and suggests mild kurtosis if the distribution were symmetric. However, kurtosis is also influenced by the skewness parameter. This parameter proves that the implied returns deviate largely from normality, with average gamma equal to 0.69. This value is far below 1 (consistently for nearly all option series) and thus the distribution far from symmetric. The shape of the implied volatility curves supports the observed skewness in the implied distributions: implied volatility curves are all downward sloping in moneyness and delta. The combination of low skewness parameter (gamma) and moderately high degrees of freedom results in extreme skewness (average -2.69) and relatively high kurtosis (average 10.50).

The left-skewness in the implied distributions far exceeds levels in actual return data (which hardly fall below -0.5) and is a result of out-of-the-money put options (and in-the-money call options) being relatively expensive. Apparently, risk-premia (the price for protection against large negative shocks in the underlying index) cannot be neglected in the price formation of options and the assumptions of continuous and costless hedging opportunities do not even hold in this very liquid market of FTSE-100 index futures and options. Since parameter estimates are quite stable over the years (Table 7.3) and across maturity (Table 7.4), risk premia play a role in all option prices. Changes in the risk-neutral distribution may therefore be attributed to changes in market expectations as well as changes in risk aversion of representative agents.

Year	# obs	σ	γ	$\ln(\alpha)$
1995	176	13.9% (1.8%)	0.79 (0.07)	3.10 (1.89)
1996	184	12.7% (2.0%)	0.73 (0.08)	2.21 (1.17)
1997	192	18.7% (6.4%)	0.70 (0.08)	2.51 (1.62)
1998	197	27.7% (7.7%)	0.64 (0.09)	2.76 (2.17)
1999	187	24.9% (5.5%)	0.60 (0.10)	2.56 (1.83)
Total	936	19.7% (8.0%)	0.69 (0.11)	2.63 (1.79)

Table 7.3 *Parameter estimates skewed Student-t method per year*

This table presents the average parameter estimates (standard deviation in parentheses) of the implied risk-neutral distribution with the skewed Student-t method. The parameters are obtained with call option data on the FTSE-100 index options in the period 1995 till 1999. The parameter σ measures the volatility of the distribution, the parameter γ determines skewness; $\ln(\alpha)$ is the natural logarithm of the degrees of freedom and predominantly determines the level of kurtosis. The last row contains the number of option series (a certain maturity on a certain date) per year.

Maturity	# obs	σ	γ	$\ln(\alpha)$
< 20 days	177	18.5% (6.1%)	0.74 (0.18)	2.39 (1.34)
20-40 days	243	19.1% (5.3%)	0.69 (0.13)	2.64 (1.04)
40-60 days	219	20.2% (5.0%)	0.68 (0.11)	2.50 (0.99)
> 60 days	297	20.6% (4.6%)	0.66 (0.10)	2.85 (1.19)
Total	936	19.7% (8.0%)	0.69 (0.15)	2.63 (1.18)

Table 7.4 *Parameter estimates skewed Student-t method per maturity bucket*

This table presents the average parameter estimates per maturity bucket (in trading days) of the implied risk-neutral distribution with the skewed Student-t method. The parameters are obtained with call option data on the FTSE-100 index options in the period 1995 till 1999. The parameter σ measures the volatility of the distribution, the parameter γ determines skewness; $\ln(\alpha)$ is the natural logarithm of the degrees of freedom and determines the level of kurtosis. The last row contains the number of option series (a certain maturity on a certain date) per maturity bucket.

The first alternative to the skewed-t under consideration is the simple lognormal model, where implied volatilities are the same for each exercise price. Two other alternatives are based on Shimko's method. We estimate one set of parameters by regressing implied volatility on a quadratic function of moneyness, and another set on a quadratic function of the option's delta. The final alternative is the implied trinomial tree approach, where we divide all times-to-maturity in 30 time steps. Parameter estimates for the lognormal, skewed Student-t and trinomial tree are obtained by minimizing the root mean squared error (RMSE) in option prices. For the Shimko methods we calculate an RMSE as well, by inserting the model volatility into the Black-Scholes formula.

In Tables 7.5 and 7.6 we observe that all methods clearly improve upon the standard (log)normal model. A simple quadratic function (Shimko), instead of a straight line, for the implied volatility causes an enormous decline in estimation errors. Although the conceptual difference is minimal between the two Shimko methods, the results indicate it is better to regress implied volatilities on option deltas than on moneyness. The former outperforms all other methods in each year (except 1995) as well as each maturity bucket. As discussed earlier, this has the additional advantage that no problems will be encountered outside the range of traded options, since delta lies in between zero and one.

If we look at the two more complex approaches, skewed-t and trinomial tree, we observe that the skewed Student-t approach is by far preferable above the trinomial tree in each maturity bucket and in each year. However, both methods consistently underperform compared to the delta curve-fitting method, although the skewed-t is slightly better than the original strike curve-fitting method. We checked that this underperformance is not a result of outliers. The skewed Student-t yields a better in-sample fit in only 20.5% of the option series, but the implied trinomial tree has an even lower score of 0.4%. So the simple quadratic curve-fitting method on option deltas works best, and this effect strengthens over time: the Shimko (delta) method yields a rather stable fit over the years, but the skewed-t deteriorates considerably. This result is surprising at first, since the Shimko method even contains one parameter less than the skewed-t and trinomial tree and is thus not a priori more flexible. An explanation might be that the two more complex methods do not have the appropriate flexibility to capture the market prices (especially the implied trinomial tree). Another, though somewhat related explanation is that option traders have been increasingly using Shimko(-related) methods to price options. The existence of implied volatility

skews and smiles are common knowledge since the early nineties and it is known that many option traders use skew- and smile related ad hoc rules to price options with different moneyness. These ad hoc rules are often based on a certain number of volatility basis points above or below the at-the-money implied volatility for each change in moneyness or delta. In other words: option traders in majority monitor the implied volatility curve and interpolate and extrapolate volatilities to price options on different points along the curve. The best a posteriori fit is then naturally obtained by mimicking the trader's practice. That may explain why methods that model the distribution will have a hard job to beat an implied curve fitting method.

Year	Normal	Shimko (Strike)	Shimko (Delta)	ITT	Skewed-t
1995	4.27	0.50	0.33	1.09	0.33
1996	5.78	0.66	0.38	1.59	0.41
1997	11.22	0.70	0.29	2.56	0.58
1998	21.24	1.39	0.48	4.08	1.28
1999	23.80	1.31	0.47	5.00	1.19
Total	13.47	0.92	0.39	2.90	0.77

Table 7.5 *Empirical fit per year*

This table presents the empirical fit of the skewed Student-t method and four alternatives. Fit is measured by the root mean squared error (RMSE) between actual and model call option prices. The fit is analyzed with data on the FTSE-100 index options in the period 1995 till 1999. 'Normal' refers to the lognormal implied distribution, 'Shimko (Strike)' to Shimko's method regressed on rescaled strike prices, 'Shimko (Delta)' to Shimko's method regressed on option deltas, 'ITT' to Nagot and Trommsdorffs (1999) trinomial tree with 30 time-steps, and 'Skewed-t' to our method.

Maturity	Normal	Shimko (Strike)	Shimko (Delta)	ITT	Skewed-t
<= 20 days	7.62	0.56	0.34	1.74	0.48
20-40 days	6.26	0.49	0.22	1.50	0.70
40-60 days	4.85	0.31	0.13	1.05	0.71
> 60 days	18.62	1.22	0.44	3.73	1.04
Total	13.47	0.92	0.39	2.90	0.77

Table 7.6 *Empirical fit per maturity bucket*

This table presents the empirical fit of the skewed Student-t method and four alternatives for different maturities. See Table 7.5 for an explanation.

7.3 Concluding remarks

The shape of an implied distribution gives important information that can be used for pricing other options, for comparing options on different assets and for closely monitoring changes in the markets perception of the underlying price process. In this chapter we proposed a new method to infer the risk-neutral distribution from option prices. Its main strength is its relative simplicity with separate parameters that directly capture the levels of skewness and kurtosis.

In an application to several years of FTSE-100 index options we compare the in-sample performance of the skewed-t method with the normal method (constant volatility), two implied volatility curve-fitting methods and a trinomial tree. Although all methods clearly improve upon the normal method, the curve-fitting method that regresses implied volatility on option delta outperforms the trinomial tree and skewed-t methods. Average root mean squared errors between actual and fitted prices are lower and this effect strengthens over time. The additional advantage of this method compared to the original Shimko method (that regresses volatility on moneyness or strike) is that it yields sensible bounded implied volatilities for any strike price. We therefore conclude that a curve-fitting method is preferred to price European-style options outside the available trading range.

Even though their fit is outperformed by a curve-fitting method, the two methods that focus on modeling the distribution of asset returns do have strong appeals. For the pricing of American-style options, backward valuation can only be applied in a tree, and the trinomial tree is a reasonably flexible candidate. The skewed-t method on the other hand has the appeal that its parameters relate directly to the moments of the distribution. This helps to judge whether the observed market prices are realistic or not. The volatility curve may for example look very smooth, but it might imply unrealistic levels of skewness and kurtosis, that are easily detected with the skewed-t method. Moreover, the parameters for skewness and kurtosis summarize changes in market expectations and risk awareness in a simple and direct manner.

8 Implied GARCH volatility forecasting³⁸

Repeated turbulence in financial markets incessantly reminds us that uncertainty is an unavoidable aspect of financial markets. With good reason the assessment of uncertainty, often proxied by price volatility, plays a prominent role in all areas of finance, ranging from investment decision-making, security valuation and risk management to monetary policy making. Even though the measurement and forecasting of volatility has attracted the interest of many researchers and practitioners, it remains a challenging statistical problem. Not only is it a problem of having a proper volatility model, but also of having a robust volatility forecasting method. The available models, such as GARCH or stochastic volatility, based on historical returns seem to work quite well in sample but generally perform poorly out-of-sample (Akgiray (1989), Dimson and Marsh (1990), Nelson (1992), Nelson and Foster (1995), Franses and Van Dijk (1995), and Brailsford and Faff (1996)).

In recent years there is some support for the hypothesis that the information provided by implied volatilities from daily option prices is more relevant in forecasting volatility than the volatility information provided by historical returns (Day and Lewis (1993), Jorion (1995), Christensen and Prabhala (1998), Fleming (1998), and Blair et al. (2001)). If option markets are efficient, option prices contain information about the price process of the underlying asset over the lifetime of the option. For the purpose of volatility forecasting this comes down to extracting the expectation of market participants about the development of future volatility. Therefore, the volatility estimate derived from option prices is a forward-looking (risk-neutral) estimate and eventually different from an estimate based on historical return data. It eliminates the choice of a particular historical sample period, which may result in better volatility forecasts.

Although the idea of option-implied volatility estimates is relatively simple, there is not one straightforward method to extract the information. Every proposed method relies on a number of assumptions regarding the model underlying option

³⁸ This chapter is based on: C. de Jong and T. Lehnert, 2002, "Implied GARCH volatility forecasting", ERIM research paper.

prices. In line with the large number of option pricing models, academics and practitioners have applied a multitude of methods to extract option-implied volatilities. Before the crash in 1987 the Black and Scholes (1973) model was applied mechanically in its original form; there was a nearly constant relationship between the implied volatility and the exercise price of the option. After the crash traders seem to adjust the volatility for moneyness and maturity before being plugged into the Black-Scholes model. The result is an implied volatility pattern that cannot be fully explained by the recent option pricing literature using historical returns of the underlying.

It is known that implied volatility covaries with realized volatility, but the major difficulty is to back out volatility information of the underlying from observed option prices. Since the assumptions of the Black-Scholes option-pricing framework are usually violated, it is a challenge to select the appropriate implied volatility. Previous studies try to explore information from (Black-Scholes) implied volatilities of traded options to estimate and forecast future volatility without explicitly modeling the underlying return process. A number of researchers have been extensively investigating the optimal weighting scheme for the different implied volatilities at different strikes³⁹. A method often applied is VIX, which is an S&P100 volatility index that combines a number of close-to-at-the-money implied volatilities into a single estimate (see Fleming et al. (1995) for a description). The index is constructed in such a way that it represents the implied volatility of a hypothetical at-the-money option with 22 days to maturity. It is therefore unable to capture any smile or term structure effects of the whole implied volatility surface. Those methods that adhere to the Black-Scholes implied volatility framework are not only arbitrary and theoretically questionable, but also the result often turns out to be a severely biased predictor of future volatility (Canina and Figlewski (1993), Fleming (1998) and Blair et al. (2001)).

In this chapter we present a new method to make volatility forecasts, which is based on a relatively recent set of option-pricing models that applies the time-series GARCH-methodology to option pricing (see Duan (1995), Kallsen and Taquq (1998), Ritchken and Trevor (1999), Bauwens and Lubrano (2002) and Heston and Nandi (2000)). We construct the expected future price process by deriving the relevant parameters of the GARCH option-pricing model from prices of traded options with

³⁹ See Bates (1996) for a review.

different strike and maturity. In contrast to methods for estimating and forecasting volatility that use past index returns, this method derives anticipated parameters of a GARCH process and therefore market expectations about the future price process. Since we need a forward-looking estimate in most finance applications of volatility, the characteristics of the future price process should be more informative than the ones of the historical price process. For example, volatility forecasts are popular among central bankers and other monetary policy makers for closely monitoring changes in the markets perception on interest rates and exchange rates (e.g. Federal Reserve Bank of Atlanta (1995), Deutsche Bundesbank (1995), Campa and Chang (1995), Leahy and Thomas (1996), Malz (1996 and 1997), Campa, Chang and Reider (1997)). Volatility forecasts of stock market prices are applied by an even wider number of financial players to investment decision-making, security valuation and risk management.

There are now several GARCH option pricing models available in the literature, but a very flexible one is the GARCH option-pricing model of Duan (1995). It has shown some empirical success and it is appropriate for our study. We estimate the parameters of the model by minimizing the relative pricing error between the market prices and the theoretical option prices of the FTSE-100 and DAX index options. The FTSE data covers the period January 1995 till July 2000, whereas the DAX data covers the period January 2000 till August 2001. Once we have the GARCH parameter estimates we can use Monte Carlo simulations to make volatility forecasts some periods ahead. As a comparison, we construct a second volatility estimate using the same GARCH specification, but now calibrated with historical returns. In an out-of-sample analysis we compare our 'implied GARCH' model with the 'historical GARCH' and conclude which method is superior in making one-day ahead forecasts of the volatility of a market index. In line with recent literature on volatility measurement we use intraday data to calculate the daily realized volatility that serves as a benchmark for our forecasts.

With this study we make several contributions to the existing volatility modeling and forecasting literature. First, we use a new method based on the informational content of option prices. With the ever-increasing trading volumes in derivative markets, we believe the information in derivatives will become the standard for making volatility forecasts. Second, we use the information contained in the whole implied volatility surface, both across maturity and across strike price. Third, we use data that have received relatively limited attention by researchers but have become

increasingly large in trading volume. The limited attention for part of our dataset, especially the DAX index options, has the advantage that it is relatively independent of the previously done volatility research that was mainly directed to S&P100 index options.

In the next section we set up the econometric framework; in Section 8.2 we describe the data; Section 8.3 provides empirical results and Section 8.4 concludes.

8.1 Methodology

In this chapter, we focus on the GARCH option-pricing framework developed in Duan (1995) and implement a model based on the exponential GARCH (EGARCH) of Nelson (1991). Apart from the well-documented GARCH effects (see Bollerslev et al. (1992)), this process is also able to model the well-known leverage effect of stock market returns (Nelson, 1991). Volatility tends to rise in response to bad news (negative excess returns) and to fall in response to good news (positive excess returns). The form of the EGARCH specification is comparable to the non-linear asymmetric GARCH process of Engle and Ng (1993), the GJR-GARCH model of Glosten et al. (1993) and the power GARCH of Ding et al. (1993). Those studies show it is crucial to include the asymmetric term in financial time series models, because volatility shocks following negative returns are significantly larger than shocks following positive returns.

The EGARCH option-pricing model assumes the risk-neutral valuation principle, and the standard Black and Scholes (1973) model can be derived as a special case. Kallsen and Taqqu (1998) develop a continuous-time version of the model and show that the same pricing results can be derived via an arbitrage-free argument. Heston and Nandi (2000) develop a closed form solution of a GARCH option-pricing model. Examples of alternative option-pricing models are the bivariate diffusion model of Hull and White (1987), the jump-diffusion model of Naik and Lee (1990), the variance-gamma model of Madan and Milne (1991) and the stochastic volatility models of Stein and Stein (1991), Wiggins (1991) and Heston (1993).

Our choice of the GARCH option pricing model of Duan (1995) is motivated by its flexibility, the recent empirical successes of the model (see among others Amin and Ng (1994), Heynen et al. (1994), Duan (1996), and Ritchken and Trevor (1999)) and the emerging availability of numerical methods for this class of option pricing

models (see Hanke (1997), Duan et al. (1998), Ritchken and Trevor (1999), Duan and Simonato (1998a&b), Heston and Nandi (2000), and Duan et al. (2001)).

For derivative valuation models with a high degree of path dependency, computationally demanding Monte Carlo simulations are commonly used for valuing derivative securities. We use a simulation adjustment method, the empirical martingale simulation (EMS) of Duan and Simonato (1998a), which has been shown to substantially accelerate the convergence of Monte Carlo price estimates and to minimize the so called ‘simulation error’. The EMS reproduces the martingale property for the simulated sample, a characteristic of all derivative pricing models.

As a first step in our empirical analysis we derive the dividend-adjusted spot rate X_t . Given a discrete dividend series D and a futures price F_t , we use the equation

$$F_t = e^{r_f(T-t)}(S_t - PV(D)) = e^{r_f(T-t)}X_t \quad (1)$$

to derive the dividend-adjusted spot rate X_t . Here $T-t$ is the time-to-maturity of the future, r_f the risk-free interest rate, S_t the spot rate and $PV(D)$ denotes the value at time t of the dividends in between t and T . In a discrete-time economy the value of the dividend-adjusted index at time t can be assumed to follow the following dynamics:

$$r_t = \ln\left(\frac{X_t}{X_{t-1}}\right) = \mu_t + \sigma_t e_t \quad (2a)$$

$$e_t | \Omega_{t-1} \sim N(0,1) \quad \text{under the probability measure } P \quad (2b)$$

$$\ln(\sigma_t^2) = \omega + \alpha \ln(\sigma_{t-1}^2) + \beta(|e_{t-1}| - \gamma e_{t-1}) \quad (2c)$$

where μ_t represents the conditional mean; Ω_{t-1} is the information set in period $t-1$ and the combination of β and γ captures the leverage effect. Daily returns of financial time series may exhibit non-zero autocorrelation. One can account for this effect by specifying the conditional mean as an autoregressive process⁴⁰ or by allowing for a risk premium attached to time-varying volatility. The specification for the conditional

⁴⁰ Bauwens and Lubrano (2002) and Hafner and Herwatz (2001) discuss how this affects option prices.

mean we selected, includes the second alternative advocated by Duan (1995, 1999) and Heston and Nandi (2000):

$$\mu_t = r_f + \lambda \sigma_t - \frac{1}{2} \sigma_t^2 \quad (3)$$

where the risk premium λ is a constant parameter, and the term σ^2 gives additional control for the conditional mean. This specification completes the baseline EGARCH model that we use for the analysis. The parameter α measures the degree of mean reversion in that $\alpha=1$ implies that the variance process is integrated. We also tried alternative specifications for the volatility dynamics, but for the ‘implied’ GARCH calibration we experienced frequent violations of the covariance stationary condition and if we control for covariance stationarity the fit was sometimes extremely bad. In contrast, restriction of the mean reversion parameter α in the EGARCH specification to values below 1 did not cause notable problems in the estimation process.

Duan (1995) shows that under the Local Risk Neutral Valuation Relationship (LRNVR) the conditional variance remains unchanged, but under the pricing measure Q the conditional expectation of r_t is equal to the risk free rate r_f . Therefore, risk neutralization transforms the error term in the following way:

$$r_t = r_f - \frac{1}{2} \sigma_t^2 + \sigma_t \varepsilon_t \quad (4a)$$

$$\varepsilon_t \mid \Omega_{t-1} \sim N(0,1) \quad \text{under the risk-neutral probability measure Q} \quad (4b)$$

$$\ln(\sigma_t^2) = \omega + \alpha \ln(\sigma_{t-1}^2) + \beta (|\varepsilon_{t-1} - \lambda| - \gamma(\varepsilon_{t-1} - \lambda)) \quad (4c)$$

In the equations above ε_t is not necessarily normal, but to include the Black and Scholes model as a special case we assume that ε_t is a Gaussian random variable. The shift of the error term can be interpreted as an additional modification of the news impact curve, therefore also modifies the asymmetry in the volatility process. The long run stationary volatility level can be shown to be equal to:

$$\sigma_{\text{inf}} = \sqrt{\exp\left(\frac{\omega + \beta E[|\varepsilon - \lambda| + \gamma(\varepsilon - \lambda)]}{1 - \alpha}\right)} \quad (4d)$$

in which the expected value should be evaluated numerically.

A European call option with exercise price K and maturity T has at time t price equal to:

$$c_t = e^{-r_f(T-t)} E_t^Q[\max(X_T - K, 0) | \Omega_t] \quad (5)$$

We rely on Monte Carlo simulations to evaluate the option numerically. Given the value of the index X_t , we generate N standard normal random numbers to advance the dynamics one period ahead and then make the empirical martingale adjustment. We repeat this procedure $T-t$ times until maturity and arrive at N simulated prices X_T . We calculate each of the N option payoffs, take the average and discount them back to period t of option valuation. Using this procedure we compute the value of an option for all exercise prices and all maturities.

We calibrate the parameters of the EGARCH option-pricing model in (5) by minimizing the square root of the mean squared pricing error between the market prices and the theoretical call and put option prices. We use relative pricing errors as defined below,

$$RMSE = \sqrt{\frac{1}{2M} \min \sum_{i=1}^n \sum_{j=1}^{m_i} \left(\frac{\hat{c}_{i,j} - c_{i,j}}{c_{i,j}} \right)^2 + \left(\frac{\hat{p}_{i,j} - p_{i,j}}{p_{i,j}} \right)^2} \quad (6)$$

where $2M$ is the total number of call and put options evaluated, the subscript i refers to the n different maturities and subscript j to the m_i different strike prices in a particular maturity series i . We use relative instead of absolute pricing errors in order to give options with different levels of moneyness equal weight in the calibration process. As starting values for the calibration we use the time-series estimates from the EGARCH model using approximately three years of historical returns.

After we calibrated the time-series estimates⁴¹, we use two time-series parameter estimates for the option calibration: the long run volatility $\bar{\sigma}$ and the risk premium parameter λ . We do so, because the large number of parameters to be estimated can make the calibration process unstable: especially the joint identification of λ and γ is cumbersome, since both parameters control for asymmetry in the news impact curve. Since this volatility ($\bar{\sigma}$) is not an explicit parameter in the model but can be derived directly from the other parameter estimates and vice versa, we constrain ω to:

$$\omega = (1 - \alpha) \ln(\bar{\sigma}^2) - \beta E[\varepsilon_{t-1} - \lambda | + \gamma(\varepsilon_{t-1} - \lambda)] \quad (7)$$

Our choice for fixing the stationary volatility is slightly different from that of Duan (1996, 1999), and Heston and Nandi (2000). They perform a constrained calibration in which the parameters λ and the *local* volatility are restricted to the time-series GARCH-estimates. We derive the local volatility from option prices, because time-series models have most problems in accurately forecasting short-run volatility fluctuations, whereas option prices can reflect new information instantaneously. In return, we constrain the long-run volatility to its time-series estimate because it turned out to be very unstable if estimated from option prices and because news has a lower impact on long run than on short run volatility.

Our approach can now be summarized as follows. First, we use three years of historical returns to estimate the time-series GARCH process. Second, we use option prices to carry out a constrained calibration by restricting the risk premium parameter λ and the long run stationary volatility level $\bar{\sigma}$ to the estimates derived from the history of asset prices. The final calibration yields estimates of the parameters α , β , γ and local volatility σ_t .

Splitting the estimation of the parameters of our model in a ‘historical’ and an ‘implied’ part has some advantages: it is more likely that option prices contain information about the future, but for risk management purposes it would be misleading to ignore all the information contained in the history of asset prices. Therefore, the method readily exploits the combination of information about the time

⁴¹ In the time-series calibration with the DAX-data we restricted the parameter ω to ensure that the long run volatility level was equal to the relatively stable historical standard deviation.

series (the volatility risk premium and the long run volatility) and the information about the volatility dynamics contained in option prices.

Given the parameters from the ‘historical’ and ‘implied’ calibration we use the EGARCH model under probability measure P in (2) to derive a volatility forecast. The estimated local volatility level $\hat{\sigma}_t$ is a one-day ahead forecast, which contains information about the expectation of market participants about tomorrow’s volatility. We compare the forecasting ability of this estimate with that of the time-series EGARCH estimate, that we denote by $\hat{\sigma}_t^{TS}$.

Since actual volatility is a latent variable, we have to construct an accurate method to evaluate our ‘historical’ and ‘implied’ forecasts. Different methods have been proposed to compute ex post estimates of it. The simplest and most common one is the square of realized return over the data. However, as Andersen and Bollerslev (1998) clearly point out theoretically, this method produces very noisy estimates of the actual volatility because of the randomness in the return process. By sampling more frequently the randomness effect can be reduced. Theoretically, the realized volatility is then closer to the actual volatility. Empirically, this is confirmed in our dataset, where we use 5-minute intraday returns to construct a volatility estimate, denoted by σ_t^{real} . This estimate is much more stable than squared returns (see also Andersen et al. (2001a and 2001b)). It is based on the 5-minute log-returns of the closest-to-maturity index future. We use a 5-minute interval, because that yields a relatively large number of returns per day without notable bid-ask bounce problems, and because it is the frequency that Andersen and Bollerslev (1998) propose. Our first index level is the opening price. All subsequent levels are the closest to each 5-minute mark, ending with the closing price.

There is some controversy about whether or not to include the overnight return. Overnight returns are relatively large in magnitude compared to the 5-minute returns and may disproportionately impact realized volatility. On the other hand, excluding overnight returns may yield a downwardly biased estimate. Since we do not want our definition of realized volatility affect our results, we decided to compare all forecasts to realized returns both including (as do Blair et al. (2001)) and excluding the return over the previous night.

8.2 Data

We use European-style options and futures price data from the two most liquid European derivatives markets: the Frankfurt-based EUREX European Futures and Options Exchange and the London International Financial Futures Exchange (Liffe). A fully electronic exchange, EUREX was formally established in 1998 following the merger of DTB Deutsche Terminbörse (German Options and Futures Exchange) and SOFFEX (Swiss Options and Financial Futures Exchange). In total number of transactions EUREX claims to be the most liquid derivatives bourse in Europe, but in the value of the trading volume Liffe is leading. Liffe's successful trading system Liffe-Connect is largely responsible for this success and one of the motives for the take-over by Euronext in 2001. The two main indices of the German and UK stock market are the DAX-30 and the FTSE-100 respectively. The options and futures on these indices are the subject of our research and traded on the EUREX and Liffe exchanges. Below we give a description of the data, starting with the DAX-30.

We use closing prices of DAX-30 index options and transaction prices of DAX-30 futures for a period from January 2000 until August 2001. For the time-series analysis we need daily index levels for a period from January 1997 until August 2001. The raw data set is directly obtained from EUREX European Futures and Options Exchange. For index options the expiration months are the three nearest calendar months, the three following months within the cycle March, June, September and December, as well as the two following months of the cycle June, December. For index futures the expiration months are the three nearest calendar months within the cycle March, June, September and December. The last trading day is the third Friday of the expiration month, if that is an exchange trading day; otherwise it is on the exchange-trading day immediately prior to that Friday.

The FTSE-100 data are similar in many respects, including the expiration months and last trading day regulations. From Liffe we purchased the daily option closing prices and intra-day transaction prices, covering the period January 1995 till July 2000. The futures transaction prices were downloaded from their Internet-set; the option closing prices were delivered on CD-ROM, but can now also be accessed through the Internet. For the time-series analysis we rely on Datastream FTSE-100 index levels, including three earlier years of data as well. Figure 1 shows the daily

development of the DAX and FTSE-100 index levels in their respective analysis periods.

Before we run the 'implied' calibration, we compute the implied interest rates and implied index rates from the observed put and call option prices using the method of Shimko (1993) based on put-call parity (see Equation (12) in the previous chapter). It can be shown that the put-call parity holds sufficiently well for our data. Alternatively, for the option calculations we could have chosen index levels from futures closing prices. This method would be equally safe since both markets are closely integrated.

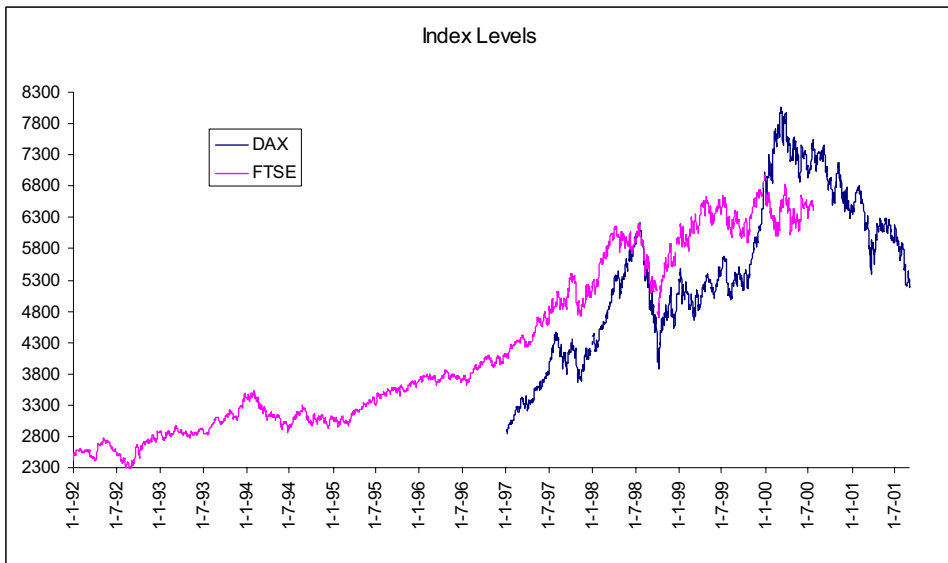


Figure 8.1 *Index levels*

The graph shows the DAX index levels for the period January 1997 until August 2001 and the FTSE-100 index levels for the period January 1992 until July 2000.

For the ‘implied’ calibration, we estimate our model using the closing prices of traded call and put options every day in our sample period. We exclude options with less than 5 and more than 75 trading days until maturity. Furthermore, instead of using a static rule and exclude options with absolute ‘moneyness’ (distance between strike and futures price) of more than 10% (see previous chapter), we use a volume rule. We exclude DAX-30 options with a daily Euro turnover of less than 10,000 Euro and a price of less than 2 Euro. In order not to lose too much FTSE-data, especially in the earlier years with relatively limited trading volume, we set the trading volume limit for the FTSE-100 options lower. We just require that at least 10 contracts of a particular option series must have been traded during the day, and furthermore exclude options with a price below £20. The choice of these particular filter rules can be motivated as follows. Since we are interested in short term volatility forecasting of the underlying index, we are interested in the information content of short term options. Second, we exclude options with less than 5 trading days to avoid liquidity-related biases. We furthermore exclude options with a price of less than 2 Euro (DAX) and less than £20 (FTSE⁴²) to avoid problems due to stale prices and problems in the minimization of relative pricing errors. Finally, we don’t automatically eliminate options whose absolute moneyness is greater than 10%, because deep in- and out-of-the-money options may still contain useful information when they are actively traded. That’s why we control for active trading by only using those options with a certain trading volume over the day. Other studies in contrast, that use a moneyness rule, run the risk of including options that are actually not actively traded and contain no information on volatilities.

Imposition of the aforementioned filtering rules reduces the average number of DAX options per trading day (puts and calls) in our dataset to around 22% of the originally more than 580 options. On average we have 124 DAX-options (puts and calls) per trading day that meet the criteria, with a minimum of 56 and a maximum of 226 options, which is by far sufficient for a reliable estimation of the four parameters. On every trading day we have at least two and at most four maturities with liquid options. On the FTSE we have fewer options per trading day, and sometimes options with just one maturity, especially in the early years of 1995 and 1996, where trading is

⁴² The lower limit of £20 may appear high, but is not really so, because average FTSE-100 option prices were £240. Options with a relatively low price receive a relatively large weight in the minimization of relative squared pricing errors. We encountered several optimization problems when we did not exclude them.

much thinner. On average there are 49 options per trading day, with a minimum of 6 and a maximum of 152. Together with the long run stationary volatility level and the risk premium parameter that we estimate from the time-series of historical returns, we ensured that a sensible estimation of the implied price process was possible.

The implied calibration is executed with the well-known Newton-Raphson or the Broyden-Fletcher-Goldfarb-Shanno procedure programmed in the statistical software package GAUSS. We use the time-series estimates as starting values in the estimation and simulate ten thousand price paths. Although this large number of simulations caused the procedure to need sometimes several minutes of computation time per trading day, it appeared to be necessary to ensure stable option values. Local optima are always possible with this type of large-scale optimizations, but convergence was nevertheless seldom a problem.

As a benchmark for our volatility forecasts we use a volatility metric based on 5-minute intra-day returns. Those returns are constructed from the contemporaneous index futures transaction prices. Every day in our sample there are several traded futures, each with a different time-to-maturity. We select the future closest to maturity. Since we are using transaction prices, negative autocorrelation may be present due to the bid-ask bounce. However, the bid-ask spreads in the two very liquid futures series are minimal, normally a fraction of a percentage. We do find some evidence for negative autocorrelation, but it is very low in magnitude and unlikely to have impacted the realized volatility estimates much. On the FTSE-100 for example, the average daily serial autocorrelation equals only -1.64% .

FTSE-100 options and futures are traded from 8:30 till 17:30 in the first part of our sample till July 17, 1999. At that date trading hours were extended to 18:00. DAX options and futures are traded from 9:00 till 17:30 in the first part of our sample (January - June 2000), and till 20:00 in the rest of the sample. The first option trade generally takes place several minutes after opening. Therefore, for the calculation of FTSE-100 realized volatility we take as the first return the difference between the opening price and the price at 8:40. Similarly, for the DAX we take the difference between the opening price and the 9:15 trading price as the first return. For the rest of the day we take the subsequent 5-minute returns. Since there are trading prices available just before and just after every 5-minute stamp, we had to make some selection. For the FTSE-100 realized volatility we take the price just following the 5-minute mark, whereas for the DAX we use the average of the prices preceding and following it. The realized volatility is computed as the square root of the sum of the

squared intra-day returns. We tested that other methods for defining the 5-minute interval return yield similar realized volatility estimates. Moreover, we constructed a second estimate of realized volatility that includes the return from the closing of previous day to the first price of that day. Obviously, overnight returns are on average quite a bit larger in absolute terms than intra-day 5-minute returns, so the realized volatility that includes the overnight return is higher on average. For the FTSE-100 the annualized volatility based on squared daily returns is on average 15.8% compared to 14.1% based on 5-minute returns and 16.2% including overnight returns. For the DAX the statistics are 18.2% based on squared daily returns, compared to 19.3% based on 5-minute returns and 21.7% including overnight returns. With the DAX-data as an example, Figure 8.2 shows that daily squared returns are a noisy estimate of volatility and using it as a benchmark for the forecasting exercise in Section 8.3 would be inappropriate (see Blair et al. (2001)).

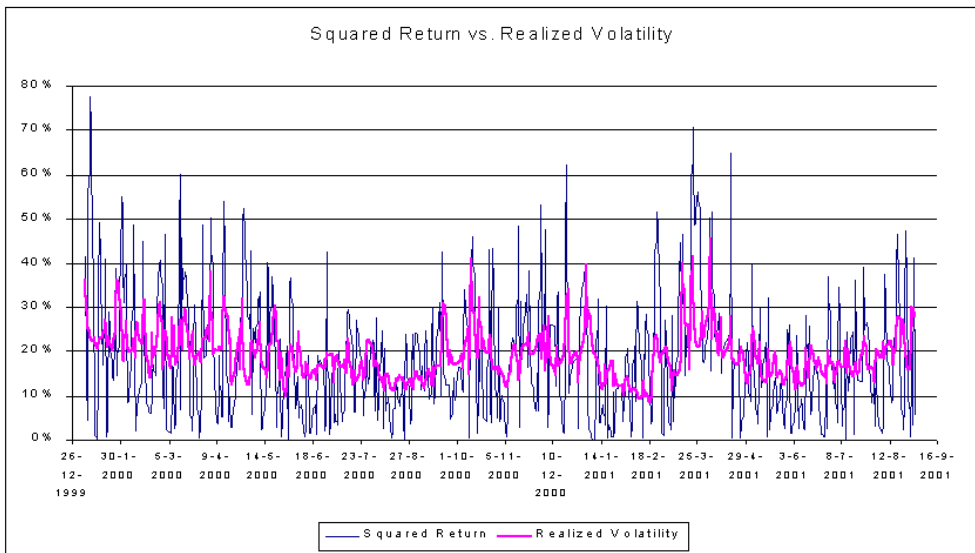


Figure 8.2 *Squared return versus realized volatility*

This figure shows the comparison of a volatility estimate based on daily squared returns and a volatility estimate based on intra-day returns. The graph is based on DAX-30 futures prices in the period January 2000 until August 2001.

8.3 Empirical results

This section describes the parameter estimates and comparisons of explanatory power between the ‘implied’ and historical’ GARCH model.

8.3.1 Parameter estimates

The ‘historical’ and ‘implied’ GARCH models are estimated daily using the implied index levels from futures prices and index option prices, respectively. As a benchmark for the in-sample fit of our GARCH option-pricing model, we use the ad-hoc Black-Scholes model of Dumas, Fleming and Whaley (1998). We allow each option to have its own Black-Scholes implied volatility depending on the exercise price K and time-to-maturity T , and use the following quadratic functional form for σ_{ij} :

$$\sigma_{ij} = \omega_0 + \omega_1 M_i + \omega_2 M_i^2 + \omega_3 T_j + \omega_4 T_j^2 + \omega_5 M_i T_j \quad (8)$$

where σ_{ij} denotes the implied volatility and M_i the moneyness⁴³ of an option with the i -th exercise price and j -th maturity. For every exercise price and maturity we compute the implied volatility and derive option prices using the Black-Scholes model.

⁴³ Moneyness is defined here as K_i/F_j , with F_j being the futures price with maturity T_j .

'Historical' GARCH			'Implied' GARCH			Ad-hoc Black-Scholes		
Parameter	Mean	(SD)	Parameter	Mean	(SD)	Parameter	Mean	(SD)
Panel A: DAX-30								
λ	0.067	(0.044)	λ	0.067	(0.044)	ω_0	1.615	(0.504)
ω	-0.409	(0.041)	ω	-0.556	(0.562)	ω_1	-2.318	(0.931)
α	0.965	(0.004)	α	0.958	(0.054)	ω_2	0.936	(0.456)
β	0.144	(0.013)	β	0.240	(0.144)	ω_3	-1.202	(1.482)
$\gamma^*\beta$	0.050	(0.008)	$\gamma^*\beta$	0.118	(0.063)	ω_4	0.735	(2.929)
						ω_5	1.007	(1.051)
RMSE				0.0786	(0.043)		0.0664	(0.034)
Panel B: FTSE-100								
λ	0.056	(0.019)	λ	0.056	(0.019)	ω_0	2.073	(2.008)
ω	-0.246	(0.087)	ω	-0.813	(1.107)	ω_1	-3.241	(3.989)
α	0.981	(0.008)	α	0.927	(0.107)	ω_2	1.344	(1.984)
β	0.083	(0.038)	β	0.142	(0.175)	ω_3	-0.724	(1.407)
$\gamma^*\beta$	0.067	(0.017)	$\gamma^*\beta$	-0.220	(0.155)	ω_4	-0.128	(0.561)
						ω_5	0.818	(1.370)
RMSE				0.0198	(0.019)		0.0126	(0.011)

Table 8.1 *Parameter estimates*

This table reports mean and standard deviations of the parameter estimates from the daily maximum likelihood or least squares estimation. Panel A contains the statistics for the DAX-30 (January 2000 – August 2001), panel B for the FTSE-100 (January 1995 – July 2000). For the implied GARCH model, ω is actually a result of setting the long run volatility equal to the long-run volatility from the time-series estimation and therefore updated but not estimated. For the implied GARCH model the risk-premium parameter λ is directly taken from the time-series estimation. Note that for the GARCH models the conditional variance (not reported) is estimated simultaneously with the other parameters in the optimization procedure. RMSE is the root mean squared error of relative pricing errors. Number of observations for the time series calibration = 753 trading days.

Table 8.1 reports the parameter estimates for the time-series GARCH, the implied GARCH and the ad hoc Black Scholes model. Note that the implied GARCH risk premium is taken from the time-series calibration. The results show that the parameter estimates vary over time, though most estimates are relatively stable. It can be shown that pre-specifying the mean in the mean-reverting volatility model stabilizes the estimation process and therefore the estimates. For the time series GARCH calibration the effect is known and the results are not surprising, but for the option implied GARCH calibration the method might be appropriate to estimate the parameters more efficiently and to avoid local maxima. The evaluation criteria for the option pricing models, the root mean squared error (RMSE) defined earlier is on average lower for the ad-hoc Black-Scholes model⁴⁴. However, the GARCH option-pricing model with the constraint that the long run volatility is equal to the sample standard deviation has a competitive fit in-sample. The results of Heston and Nandi (2000) show that the ad-hoc Black-Scholes model might achieve better in-sample fit only by overfitting the data, but underperforms GARCH option pricing models out-of-sample. An out-of-sample pricing analysis is beyond the scope of this study, but we can conclude that the pricing performance of our method is reasonably accurate. The resulting one day ahead volatility forecast produced by the ‘historical’ and ‘implied’ GARCH models are presented in Figure 8.3.

⁴⁴ This is similar to the results in Chapter 7, in which we showed that the volatility curve-fitting methods achieved a better in-sample fit than distribution-based methods.




Figure 8.3 *Option implied volatility forecast vs. time series volatility forecast*
 The graphs show the annualized one-day ahead volatility forecasts of the option implied GARCH model and the time series GARCH model. The first graph is based on the DAX-30 data in the period January 2000 until August 2001. The second graph is based on the FTSE-100 data in the period January 1995 until July 2000.

8.3.2 Lead-lag relations

It can be seen from the DAX-graph that in general both methods lead to different volatility predictions, but move fairly well in line. On the German market most of the time the ‘implied’ GARCH volatility forecast (average = 22.1%) is lower compared to the ‘historical’ GARCH forecast (average = 23.7%), but during some periods the ‘implied’ is exceeding the ‘historical’. In the UK market we observe the opposite: the average implied forecast (18.3%) exceeds the ‘historical’ forecast (16.9%). Another difference is the much lower stability of the implied forecasts on the UK market than on the German market. Especially during the financial turbulence in the second half of 1997 and 1998 we obtain largely fluctuating implied forecasts, sometimes exceeding 50%. This is however in line with the realized volatility estimates, which are high and fluctuating as well in those periods.

Visual inspection of the first graph suggests that whenever news is entering the market leading to a rising volatility estimate for the ‘historical’ GARCH model over the following days, this news is already incorporated in the ‘implied’ volatility forecast and the ‘implied’ forecast is suddenly exceeding the ‘historical’ one. Therefore, a positive or negative jump in ‘implied’ volatility forecast seems to indicate that there is new information in the market, but the ‘historical’ GARCH model needs some days to update the volatility estimate. As a result the time series of both local volatility estimates suggests that the ‘implied’ forecast is leading the ‘historical’ forecast on the German market, but possibly not on the UK market. We want to test this hypothesis by conducting a Granger causality test (Granger (1969)). The method determines the causal directions between two variables by indicating if changes in one variable induce changes in the other variable or if both variables are jointly determined. Under the hypothesis of one variable not Granger-causing the other variable, the test statistic has the F-distribution and a rejection of the null hypothesis indicates causality. Table 8.2 reports the results for the Granger causality test.

Null Hypothesis	Lag = 1		Lag = 3		Lag = 5	
	F-stat	Prob.	F-stat	Prob.	F-stat	Prob.
Panel A: DAX-30						
'Historical' forecast does not Granger cause 'implied' forecast	1.22	0.27	1.81	0.15	2.20	0.05
'Implied' forecast does not Granger cause 'historical' forecast	96.30	0.00	56.20	0.00	33.40	0.00
Panel B: FTSE-100						
'Historical' forecast does not Granger cause 'implied' forecast	12.85	0.00	0.92	0.43	0.88	0.49
'Implied' forecast does not Granger cause 'historical' forecast	2.85	0.09	1.54	0.20	0.99	0.42

Table 8.2 *Granger causality test*

This table presents the results of a Granger causality test for 'historical' and 'implied' volatility forecasts. Panel A contains the statistics for the DAX-30 (January 2000 – August 2001, 424 observations), panel B for the FTSE-100 (January 1995 – July 2000, 1394 observations). 'Lag' is the number of lags used in the regression. 'F-stat' is the Wald F-statistic and 'Prob.' is the corresponding p-value. The p-values for the DAX-30 indicate that the 'implied' is leading the 'historical', but no such relation can be established for the FTSE-100.

The test results on the UK market have somewhat contradictory outcomes for the different lag lengths. It seems most safe to interpret this as that no series leads the other. Since the FTSE-100 covers a relatively long history with much lower liquidity in the early years, we performed the Granger causality test also on (roughly) the second half of our dataset, from January 1998 onwards. Then results point weakly to the implied leading the historical forecast. We can reject with p-values of 7%, 6% and 20% the second hypothesis, while we can only reject the first hypothesis at higher p-values of 8%, 17% and 74%. The results are not convincing enough to draw any definitive conclusions, but we will explore this issue further when we formally analyze the predictive power of the different volatilities.

On the German market the results are clearer. The test strongly suggests that our conjecture about a causality in the German market is statistically significant: the hypothesis that the option implied forecast does not cause the time series forecast can be rejected of all reasonable significant levels, while we cannot reject the hypothesis that the ‘historical’ forecast does not cause the ‘implied’ forecast on a 10% significance level for lags equal to 1 and 3 and on a 5% significance level for a lag equal to 5 trading days.

So far we conclude that the volatility estimate derived from option prices seems the volatility estimate using historical return data on the German market, but does not seem to lead nor lag in the UK market. This would mean that the options market is more informative when forecasting DAX-30 volatility than FTSE-100 volatility, which we test below.

8.3.3 Forecasting power

In the following the out-of-sample accuracy of the volatility forecasts is compared. Given the ‘historical’ volatility forecast $x_{TS,t+1}$ and ‘implied’ volatility forecast $x_{OPT,t+1}$ made at time t of the realized volatility y_{t+1} known at time $t+1$, we can evaluate both models by comparing the multiple R^2 statistics from the regression

$$y_{t+1} = \alpha + \beta_{TS}x_{TS,t+1} + \beta_{OPT}x_{OPT,t+1} + u_{t+1} \quad (9)$$

The multiple R^2 statistics can be interpreted as a measure of information content of the mixture of forecasts, which have more predictive power than univariate forecasts (Day and Lewis (1992)). Table 8.3 reports the (multiple) squared correlation R^2 from regressions of realized volatility on one or two volatility forecasts. Table 8.4 reports F-statistics of the test that only one of the predictors forecasts equally well as both the predictors. The results are very much in line with the Granger causality test statistics.

Forecasting aim	Forecasting method	α	β_{TS}	β_{OPT}	R^2
Panel A: DAX-30					
Realized vol. excl. overnight	'Historical'	0.035 (2.4)	0.668 (11.2)		23.00%
	'Implied'	0.030 (2.9)		0.738 (16.0)	37.90%
	'Historical' + 'Implied'	0.008 (0.6)	0.194 (2.8)	0.630 (10.5)	39.00%
Realized vol. incl. overnight	'Historical'	0.013 (0.8)	0.859 (11.6)		24.20%
	'Implied'	0.002 (0.2)		0.973 (17.4)	41.90%
	'Historical' + 'Implied'	-0.023 (-1.4)	0.218 (2.6)	0.851 (11.7)	42.80%
Panel B: FTSE-100					
Realized vol. excl. overnight	'Historical'	0.026 (7.8)	0.678 (36.5)		48.80%
	'Implied'	0.059 (21.9)		0.452 (34.1)	45.50%
	'Historical' + 'Implied'	0.025 (8.0)	0.430 (17.0)	0.238 (13.7)	54.90%
Realized vol. incl. overnight	'Historical'	0.019 (3.7)	0.848 (29.6)		38.60%
	'Implied'	0.053 (13.5)		0.602 (31.0)	40.80%
	'Historical' + 'Implied'	0.017 (3.6)	0.456 (11.8)	0.375 (14.0)	46.10%

Table 8.3 *Regression results*

This table reports parameter estimates and t-statistics (in parentheses) from regressions of 'historical' and 'implied' GARCH forecasts on realized volatility, both including and excluding overnight returns. The parameter β_{TS} is the coefficient of the historical time-series GARCH volatility, β_{OPT} of the implied GARCH local volatility (see regression 9). Panel A contains the statistics for the DAX-30 (January 2000 – August 2001, 424 observations), panel B for the FTSE-100 (January 1995 – July 2000, 1394 observations). The R^2 measures the explanatory power of the different predictors.

Forecasting aim	Null hypothesis	DAX-30		FTSE-100	
		F-stat	Prob.	F-stat	Prob.
Realized vol. excl. overnight	'Historical' does not improve forecast	6.17	0.23%	143.65	0.00%
	'Implied' does not improve forecast	147.42	0.00%	86.25	0.00%
Realized vol. incl. overnight	'Historical' does not improve forecast	4.56	1.09%	91.40	0.00%
	'Implied' does not improve forecast	161.99	0.00%	136.75	0.00%

Table 8.4 *Additional Forecasting Power*

This table reports the results of an F-test to determine whether the two regressors 'historical' and 'implied' in equation (7) explain realized volatility better than only one regressor. We report the statistics for the DAX-30 (January 2000 – August 2001, 424 observations) and the FTSE-100 (January 1995 – July 2000, 1394 observations). The p-value indicates the probability of rejecting the null hypothesis.

First, it can be observed that the predictors are able to explain a large chunk of realized volatility. The squared correlation R^2 coefficients are above or in the upper range of the usually reported levels of around 10-40%. We believe this is partly due to our definition of implied volatility, but also due to the definition of realized volatility that is much more accurate than squared daily returns.

Second, we analyze whether the predictors are unbiased or not. For the predictors to be unbiased the regression coefficients should be equal or sum up to unity and the constant should be undistinguishable from zero. On the FTSE we find that none of the predictors is unbiased, which is in line with the existing research on the forecasting power of implied volatility (Fleming (1998) and Blair et al. (2001)). This does not mean that the historical or implied forecasts consistently over- or undershoot the actual volatility: their averages do not deviate much from the average realized volatility, as shown earlier. Biased in this setting means that in times of high predicted volatility actual volatility can be expected to be somewhat lower and vice versa in times of low predicted volatility. On the DAX however, our predictors are much less biased, especially if we include overnight returns. In the latter case we cannot reject the hypothesis that the implied and time-series forecasts are unbiased: the constants are close to one, and the two coefficients equal or sum up to a value close to one.

It is seen finally that on the German market (DAX-30) the ‘implied’ forecast clearly outperforms the ‘historical’ forecast. The historical GARCH method performs rather poorly and the incremental information of the implied method is striking: the value of R^2 increases by 16-18% when the realized volatility estimate is regressed on the historical and implied forecasts instead of historical forecast alone, whilst the increase is only 1% when the realized volatility estimate is regressed on the historical and implied forecasts instead of implied forecast alone. Therefore, the information from the time-series provides hardly any additional information compared to the option-implied information (see also Table 8.4). On the UK market however, both predictors perform comparably well and none of the two should be disregarded in making volatility forecasts, since a weighted combination drives the predictive ability up by 5-10% to around 50%. We suspected that this difference between the German and UK market might have something to do with the longer history of the FTSE-data, but results are qualitatively unchanged in the second half of our FTSE-dataset. An alternative explanation is that the London option market is less efficient than the market in Frankfurt and does not incorporate all the information that can be inferred from standard time-series models, even though the predictive power of the implied is relatively high. However, with most experienced European traders working in London, this does not sound very plausible and leaves us with a puzzle.

8.4 Concluding remarks

This chapter is concerned with short-term volatility forecasting. We compare and combine the information in historical returns and the information in option prices to investigate what source contains most valuable information in forecasting FTSE-100 and DAX-30 volatility in the period of January 1995 till July 2000 and January 2000 till August 2001 respectively. More particularly, we compare the forecasts of a time-series EGARCH model to the forecasts of an EGARCH model whose main parameters are derived from contemporaneous option prices. We use the Duan (1995) option-pricing model to identify an option-implied EGARCH process and the corresponding 1-day ahead volatility forecast. A large number of simulations and optimizations is required to identify the parameters of the ‘implied’ model. Those parameters vary over time, but are relatively stable and provide an accurate fit to the option prices.

Our results yield different outcomes on the two markets. In the German market (DAX-30), the implied volatility forecast is leading the historical forecast. The first forecast Granger-causes the second, but not the other way round. This is an indication that information is more quickly compounded in option prices than in the most recent returns. In the UK market (FTSE-100) to the contrary, we do not find convincing evidence for any lead-lag relation.

This lead-lag intuition is confirmed in the out-of-sample 1-day ahead volatility forecasts. We test the implied and historical forecast accuracy against a realized volatility estimate that is constructed with 5-minute intra-day returns. This is a more reliable estimate of actual volatility than the very noisy squared returns. In the German market, the historical forecast is able to explain around 23% of realized volatility, whereas the same ability of the implied estimate is 1.5 times larger. In the UK market both predictors perform equally well and explain between 39% and 49%, depending on whether we include or exclude overnight returns in constructing realized volatility. In the UK market it is wise to construct a weighted combination of the two predictors, since that explains even 8% extra, but in the German market the implied forecast alone contains nearly all information. Moreover, in the German market the implied forecast, alone or combined with the historical forecast, is a relatively unbiased predictor of realized volatility including overnight returns. We believe this is a new and important result. Our result is also new that options on the leading German index DAX-30 contain more (if not all) information than the history of prices, contrary to the UK market. Options on the main UK index FTSE-100 do not contain more valuable information than the history of prices, although both yield powerful forecasts.

Another contribution of this study is that we introduce a different and theoretically sound methodology to extract the information from the option prices. Previous studies try to explore information from Black-Scholes implied volatilities of traded options to estimate and forecast future volatility without explicitly modeling the underlying return process. In those studies various optimal weighting schemes are being proposed for the different implied volatilities at different strikes. In doing so, those methods ignore the information contained in the volatility smile pattern, and probably more importantly, in the volatility term structure. We do not compare our method to any such method, because a severe selection bias would be in place. Since different Black-Scholes implied volatilities could be combined in so many ways, the choice for one particular method would largely impact our comparison.

Apart from their theoretical drawbacks, we believe the ignorance of smile and term structure information by Black-Scholes-based methods may hurt their forecasting ability, especially at longer horizons. Therefore, in further research we will extend the forecasting to longer horizons of up to several weeks ahead. This opens up new interesting fields of research that are not solely focused on the second moment. For example, as an application of return distribution forecasting, we plan to calculate Value-at-Risk estimates and analyze their accuracy.

Another extension that we plan to make in future research is a comparison of our implied GARCH volatility forecasts with a forecast based on high-frequency data. Such a comparison seems logical given the benchmark of realized volatility (based on high-frequency data) we employ. Although such a high-frequency forecast is solely based on historical returns, it has very recently shown to yield good forecasting results (Martens and Zein (2002), Li (2002), Pong et al (2002)). High-frequency models are able to respond more quickly and accurately on market movements than GARCH-type specifications with returns that are aggregated on a daily basis. When a long memory effect is incorporated in these models, they can even compete with and sometimes outperform implied volatility models at longer horizons. This makes them a natural challenger of the implied GARCH model in future research.

9 Pricing the spikes in power options ⁴⁵

Since the early 90's electricity markets are being reformed worldwide from a highly government controlled and vertically integrated environment into competitive markets. Before the deregulation, government authorities fixed prices based on (marginal) production costs in a very predictable manner. Now that many wholesale markets are deregulated, market participants have to get used to an environment with very volatile prices and high uncertainty. Participants face the additional complexity that volatility far exceeds the volatility in markets that are considered relatively risky, such as those for stocks, bonds, and other commodities. At the same time, the number of available instruments to control risks has grown radically. Markets gradually extend trading in day-ahead physical deliveries (spot or pool), and forward contracts with physical deliveries, to more advanced physical and financial products, such as swaps, futures, options, caps, floors and spark spreads. Most derivatives trade in over-the-counter markets, but increasingly on exchanges as well. Examples of such trading venues are the NYMEX, Nord Pool, European Energy Exchange, Chicago Board of Trade, Minneapolis Grain Exchange, Sydney Futures Exchange, and New Zealand Futures Exchange.

In this chapter we focus on the valuation of options on spot⁴⁶ prices. Although options on spot (day-ahead) electricity form only a subcategory of tradable electricity contracts, their valuation is an economically important topic. First, options on spot prices are embedded in many contracts. For example, floating-price contracts with a minimum or maximum (cap or floor) contain in fact a series of call or put options, and are increasingly popular products among end users. Their tradable counterpart is a daily exercisable option, with exercise opportunities each day for a period of weeks,

⁴⁵ This chapter is partly based on: C. de Jong and R. Huisman, 2002, "Option Formulas for Mean-Reverting Power Prices with Spikes", ERIM research paper.

⁴⁶ With spot electricity prices, we mean the prices for electricity that are determined one day in advance on spot exchanges, in pool systems or over-the-counter markets. We can have prices for time periods ranging from individual quarters of an hour, hours, blocks of hours (e.g. peak and off-peak) to daily averages (baseload).

months or year⁴⁷. Finally, the valuation of options on spot electricity is important in valuing real assets. Flexible energy production capacity provides an option to produce or not in the hours or days ahead. The value of such flexible capacity is therefore equal to the value of a series of call options on the spot (short-term) price⁴⁸, and our framework provides a basis to value those assets.

For the valuation of electricity derivatives we cannot simply rely on models for financial and other commodity contracts. Electricity is a pure flow commodity with limited storability and transportability that strongly affect the behavior of electricity spot and derivatives prices. This lack of flexibility causes spot prices to depend largely on local and temporal supply and demand conditions. If demand and supply would respond promptly to price movements, prices would not deviate much from other commodity prices. The elasticity of both supply and demand however is relatively limited (see for example Pirrong and Jermakyan, 2000). Only a few large industrial customers have the flexibility to vary their power demand in response to market conditions, whereas most power plants can gear up generation only with a significant time lag. This time lag causes occasional extreme prices, called spikes, which revert within hours or days to a more stable level. All this results in the well-documented characteristics of electricity spot prices, including spikes, mean-reversion, large seasonal variations and extremely high volatility.

These peculiar characteristics of electricity prices have induced researchers to develop special models for electricity prices. Such models are the basis for risk management applications, for the pricing of physical and financial contracts, and for the valuation of real assets. With the increasing number of tradable contracts, the main challenge for researchers is the development of models to price those contracts. We recognize two different sets of electricity contract valuation approaches. The first approach is most popular among academics and consists of modeling simultaneously spot and forward contracts. Examples are Schwartz (1997), Hilliard and Reis (1998), Pilipovic (1998), Pirrong and Jermakyan (1999, 2000), Deng (2000), and Lucia and

⁴⁷ Daily exercisable options are the only tradable spot option contracts (to our knowledge). They can either be for physical delivery or financially settled. Typically, the holder of the option needs to indicate in the morning before the delivery day whether to exercise or not. This may be before or after settlement of the spot market. In the latter case, suboptimal exercise is possible.

⁴⁸ If fuel costs are volatile as well, then a generation asset can be considered an option on the difference between the electricity price and fuel costs, the spark spread.

Schwartz (2002). This approach faces the difficulty that standard arbitrage principles cannot be applied to map spot prices to forwards and futures. Therefore, the proposed solutions are naturally derived from the bond or storable-commodity pricing literature. Risk-neutral processes are obtained either through the specification of risk premia or convenience yields. A theoretical drawback of models based on convenience yields is that electricity is not storable, and therefore the interpretation of convenience yields is questionable. Moreover, fitting the theoretical forward curve to market data is a serious problem, since data is limited and several institutional factors influence power forward price dynamics.

The second approach describes how to price options on spot, forwards or futures, and takes the forward curve as given. We take this approach that makes option valuation consistent with market prices, and we believe this is necessary to bridge the gap between academic theory and the derivative models that are predominantly being used in practice (Black & Scholes (1973) and Black (1976)). For option valuation we simply need models that adequately describe the dynamics of spot or forward prices, align them with the market forward curve, and then use arbitrage or “fair-pricing” principles to derive option prices⁴⁹. Other examples of this approach are Miltersen and Schwartz (1998), Clewlow and Strickland (1999), Bjerksund, Rasmussen and Stensland (2000) and Koekebakker and Ollmar (2001).

Price returns of longer-term futures and forwards fulfill the conditions for normality relatively well, and hedging related options with forwards or futures is often feasible. Consequently, standard arbitrage-based pricing-formulas may yield reliable results for options on longer-maturity forwards and futures. Spot returns however are clearly not lognormally distributed (see e.g. Lucia and Schwartz, 2002) and the standard option pricing formulas may yield totally incorrect outcomes. That’s why we need a different type of pricing approach, especially for further in-the-money and out-of-the-money options.

This chapter presents such an approach for determining the ‘fair value’⁵⁰ of options on electricity spot prices. It is similar in spirit to Clewlow and Strickland

⁴⁹ It should be noted that forward trading is often liquid for only a few maturity series, but combinations of market prices with bottom-up models (see e.g. Fleten and Lemming (2001)) may be used to complete the forward curve.

⁵⁰ Since spot price risk cannot be hedged (at least not financially), no arbitrage-free pricing results can be obtained. The pricing approach thus yields ‘fair values’ instead of ‘arbitrage-free values’.

(1999), and Lucia and Schwartz (2002) but extends their mean-reverting framework with the necessary spikes, modeled as a separate and independent regime⁵¹. At the same time it keeps the attractive feature of closed-form formulas, where other spot price models (Bhanot (2000), Deng (2000), Knittel and Roberts (2001), Huisman and Mahieu (2001), and Escribano, Pena and Villaplana (2002)) would require extensive simulations. The two regimes (one for the ‘normal’ process, one for the spikes) capture the systematic alternations between stable and unstable states of demand and supply. The price process takes the main dynamics of electricity prices into account, such as seasonality, mean-reversion and, most importantly, spikes. Furthermore, a major contribution of the model is that it allows for multiple consecutive spikes, which is important for risk management and derivative pricing purposes. Based on this spot price model we present closed-form formulas to price European-style options on spot electricity prices. We furthermore show how the underlying spot price model can be aligned with the observed forward curve in the market, which makes the option valuation consistent with market data.

The chapter is built up as follows. First, we present a spot price model that incorporates the most prominent features of electricity spot prices: mean-reversion and spikes. Next, in Section 9.2 we estimate the model parameters for Dutch APX baseload, peakload and off-peak spot prices. Section 9.3 first describes how the model can be aligned with market forward curves in a practical manner that avoids the separate modeling of seasonalities and risk premia. Then it presents closed-form formulas for European-style options and describes how other types of options can be priced. We end with some concluding remarks.

9.1 The two regimes model for spot electricity prices

A standard mean-reverting specification is relatively successful in modeling commodities such as oil and gas⁵², but not in modeling electricity, due to the existence of spikes. This ‘spiky’ behavior of electricity prices has mainly to do with the non-

⁵¹ Deng (2000) also proposes to model spikes in a regime-switching model. He derives formulas for pricing futures, forwards and standard options, but requires simulations to evaluate the outcomes.

⁵² See for example Pindyck (1999).

storability of the commodity and the relative inelasticity of demand and supply, as we discussed in the introduction. Parameter calibration generally leads to unrealistically high volatility, incorrect mean reversion parameters and too high levels to which the spot prices would converge.

Modeling spikes in a satisfactory framework has turned out to be a major challenge for researchers and practitioners in electricity markets. The most common approach is the addition of a jump diffusion process to the mean-reverting process. Stochastic jump models allow for sudden extreme returns that lead to long-term shifts in price levels. They are quite successful in stock markets, but do not incorporate an important characteristic of electricity prices: spikes are relatively short-lived. A jump diffusion process model allows for large price movements, but does not deal well with the fact that after a spike prices quickly bounce back to normal levels. In a mean-reverting jump diffusion process this can only be achieved by an unrealistically high mean reversion parameter that forces prices back to normal levels after a spike.

9.1.1 Regime-switch models

Regime-switch models have the potential to solve some of these deficiencies, since they allow for distinct time-series behavior in different periods of time. The basic regime model has the following specification (Hamilton, 1989):

$$\ln S_t = \mu_{\lambda(t)} + \varepsilon_t \quad \text{where} \quad \varepsilon_t \sim N(0, \sigma_{\lambda(t)}^2) \quad (1)$$

Here $\lambda(t)$ is a latent variable representing the regime of the process in time period t . The process can thus be in one of the regimes at each time t . Huisman and Mahieu (2001) propose a regime-switch model with three regimes: there is a mean-reverting regime with moderate mean-reversion and volatility, an initial jump regime that models the process when prices suddenly increase or decrease, and a subsequent jump regime, that describes how prices are forced back to the stable regime. The two jump regimes both have a more extreme expected return and volatility than the mean-reverting regime. The subsequent jump regime has a zero probability of occurrence if prices in the previous time period are in the mean-reverting regime, but a probability of one if they are in the initial jump-regime. The main drawback of this model is that it does not practically allow for multiple consecutive jumps, which are a frequent

phenomenon in electricity markets (see Figure 9.2) and crucial for risk management purposes and derivative valuation. Therefore, we introduce a model with only two regimes: a stable mean-reverting regime, and a spike regime. It might seem surprising that the omission of one regime gives the model the flexibility to capture consecutive jumps. However, we don't need a third regime to pull prices back to stable levels, because we assume that prices in the two regimes are independent from each other. Put differently, if there is a generator outage for example, prices may be high for some time period, but once the generator is repaired, prices continue as normal. We believe this regime specification fits well with the structure of electricity markets and is confirmed in the data. As a side effect of the independence of the two processes, we can combine closed-form formulas of mean-reverting time series, with those of lognormally distributed spikes to simplify derivative valuations and to circumvent time-consuming Monte Carlo simulations.

9.1.2 The two-regime framework

The first step in modeling electricity spot prices, P_t , is the separation of the predictable component from the stochastic component (Hamilton, 1994).

$$p_t = \ln P_t = f(t) + x_t \quad (2)$$

The first component, $f(t)$ accounts for predictable regularities, such as any genuine periodic behavior and any trend, and is a deterministic function of time. Seasonalities can be modeled with for example sinusoidal functions or with dummy's for different seasons, to which we come back in Section 9.2. The stochastic second component, x_t , is the more interesting and we continue with its specification below. In the remaining we refer to the stochastic part x_t as the "log spot price", but remember that in fact it is the log spot price from which predictable trends are removed.

In the two-regime framework, we assume that the spot price of electricity can be in one out of two regimes at each time period t . The first regime reflects the normal behavior of electricity prices and the second reflects dynamics in case of spikes. We assume that the deterministic trend $f(\cdot)$ in (2) remains the same across regimes, since

spike data is too limited to warrant a separate seasonal specification. We then specify the two-regime model as follows:

$$p_t = f(t) + x_{\lambda(t),t} \quad (3)$$

where $t = 1, \dots, T$ and $\lambda(t) = M, S$. The latent variable $\lambda(t)$ may assume two values. We refer with $\lambda(t) = M$ to the mean-reverting regime and $\lambda(t) = S$ to the spike regime. So, $x_{M,t}$ is the stochastic process for the mean-reverting regime and $x_{S,t}$ is the stochastic process for the spike regime. For the first regime we specify a standard mean-reverting process.

$$\text{Mean-reverting regime: } x_{M,t} = x_{M,t-1} + \alpha(\mu_M - x_{M,t-1}) + \varepsilon_{M,t} \quad (4)$$

$$\text{where } \varepsilon_{M,t} \stackrel{i.i.d.}{\sim} N(0, \sigma_M^2)$$

The parameter μ_M is the long-run stationary level for the natural logarithm of spot prices. It determines to what value spot prices converge. The parameter α measures the speed of convergence from the current to the long-run level and is related to the concept of half-life, a well concepts in physics: the time it takes to move on average halfway from the current level to the long-term level⁵³. The spikes of the second regime are modeled with a simple lognormal distribution whose standard deviation and mean⁵⁴ are higher than those of the mean-reverting process. We have the following specification:

$$\text{Spike regime: } x_{S,t} = \mu_S + \varepsilon_{S,t} \quad (5)$$

$$\text{where } \varepsilon_{S,t} \stackrel{i.i.d.}{\sim} N(0, \sigma_S^2)$$

⁵³ Half-life = $\ln(0.5)/(1-\alpha)$

⁵⁴ So-called off-peak hours (when demand is low) are characterized by negative spikes, due to the abundance of supply relative to demand during those periods of the day. In those cases the mean of the spike becomes negative.

At any point in time the price process is either in regime M or regime S. However, the model specification does not end with the two regimes, since we haven't defined the transition process yet. For this we use a Markov transition matrix, which contains the probabilities of switching from one regime to the other. With two regimes, the Markov transition matrix Π is a 2x2 matrix. The element in column j and row i contains the probability π_{ij} of going from regime i in period t to regime j in period $t+1$ ($i,j = M,S$).

$$\Pi = \begin{bmatrix} 1 - \pi_{MS} & \pi_{SM} \\ \pi_{MS} & 1 - \pi_{SM} \end{bmatrix} \quad (6)$$

If spot prices are in the mean-reverting regime today, we know that with probability π_{MS} the next day is a spike, and with probability $1 - \pi_{MS}$ the mean-reverting regime continues. Similarly, we know that a spike is followed by another spike with probability $1 - \pi_{SM}$, and otherwise the mean-reverting regime resumes.

We stated earlier that the two regimes are independent, which holds true for the prices in each regime. The above probability structure however ensures that there is a relation between the two regimes in terms of the probability that they occur. For example, when we observe a spike today, then we expect a spike tomorrow with a larger probability than when prices were normal. This is the type of relation we observe in electricity markets, but does not prevent us from disentangling option prices into two components, as long as prices are independent.

9.1.3 Parameter estimation

The parameters of the two regimes can be calibrated by maximum likelihood when we condition on the regimes. Based on the normal distribution of the error terms, the loglikelihoods have the following form:

$$LL_{M,t} = -\frac{(x_{M,t} - x_{M,t-1} - \alpha(\mu_M - x_{M,t-1}))^2}{2\sigma_M^2} - \ln \sigma_M - \frac{1}{2} \ln 2\pi \quad (7)$$

$$LL_{S,t} = -\frac{(x_{S,t} - \mu_S)^2}{2\sigma_S^2} - \ln \sigma_S - \frac{1}{2} \ln 2\pi \quad (8)$$

The two-regime specification for spot prices introduces a technical complexity in the calculation of the mean-reverting loglikelihood: the Equation (7) depends on the mean-reverting price of the previous period, which does not exist if the previous period was a spike. This means that if prices were in a spike yesterday, we do not know from what level they have to revert today (if today is a ‘normal regime’ period). We solve this issue as follows. First, we rewrite the loglikelihood function (7) in a more general form:

$$LL_{M,t}(x_{t-i}) = -\frac{(x_{M,t} - E_{t-i}[x_{M,t}])^2}{2Var_{t-i}[x_{M,t}]} - \frac{1}{2} \ln Var_{t-i}[x_{M,t}] - \frac{1}{2} \ln 2\pi \quad (9)$$

Conditional on information about $x_{M,t-i}$, the above equation gives the loglikelihood for an observed price $x_{M,t}$ in the mean-reverting regime. We now use the model’s posterior probabilities to calculate for each of the k last prevailing log spot prices (x_{t-1} , ..., x_{t-k}) the probability that it was the last (observed) mean-reverting price. Next, we calculate the loglikelihood in Equation (9) assuming alternatively that the price in period $t-1$, ..., $t-k$ was the last mean-reverting price. If we look i periods back, in the likelihood equation we use $E_{t-i}[x_{M,t}]$ and the appropriate (higher) variance $Var_{t-i}[x_{M,t}]$ to capture the effect that prices are stochastically mean-reverting. These expected values and variance can be determined recursively as follows:

$$E_{t-1}[x_{M,t}] = \alpha\mu_M + (1-\alpha) \cdot E_{t-2}[x_{M,t-1}] \quad (10)$$

$$Var_{t-1}[x_{M,t}] = (1 + (1-\alpha)^2) \cdot Var_{t-2}[x_{M,t-1}] \quad (11)$$

and the loglikelihood of the mean-reverting regime equals the probability weighted sum of the conditional loglikelihoods:

$$LL_{M,t} = \sum_{i=1}^k \Pr[\lambda(t-i) = M \wedge \lambda(t-j, j < i) \neq M] \cdot LL_{M,t}(x_{t-i}) \quad (12)$$

We set k equal to 5, because the sum of the posterior probabilities then approaches 1 well enough in our data. This solves the latent variable problem in the original Equation (7).

The likelihood of the whole process equals the weighted sum of the likelihoods of the two regimes. The weights are determined by each regime's prior probability. If we denote the prior probability of prices being in the mean-reverting regime at time t by $\Pr_{t-1}[\lambda(t) = M]$ and of being in the spike regime by $\Pr_{t-1}[\lambda(t) = S]$, then the likelihood function equals:

$$LL = \sum_t \Pr_{t-1}[\lambda(t) = M] \cdot LL_{M,t} + \Pr_{t-1}[\lambda(t) = S] \cdot LL_{S,t} \quad (13)$$

This completes the specification of the mean-reverting regime model with independent spikes. In the next section we evaluate the parameters estimated from Dutch day-ahead prices.

9.2 Model estimation results

This section reports the estimation results of the two-regime model presented in the previous section, as well as the Huisman-Mahieu (2001) and a standard mean-reverting model. We discuss some data issues and evaluate the parameter estimates to see whether the two-regime model picks up mean-reversion and spikes sufficiently well. We use those parameter estimates to price options on Dutch APX⁵⁵ spot prices in the next section.

⁵⁵ See www.apx.nl

9.2.1 Data

The Dutch power market was liberalized for large consumers on January 1st 2001. From that day onwards the wholesale day-ahead prices on the Amsterdam Power Exchange (APX) reflect the forces of demand and supply. Trading volume on this electronic exchange has increased steadily, notably since January 2002, when medium-sized consumers became free to choose their energy-supplier as well. By June 2002 the APX-volume represented approximately 14% of total electricity consumption. We use data from January 2nd 2001 till June 30th 2002 of baseload, peak and off-peak day-ahead prices, totaling 545 observations for each index. The exchange defines baseload prices as the equally weighted average of the 24 individual hourly prices. The peak prices are the average of hour 8 till 23 (7:00 – 23:00); the off-peak prices are the average of the 8 remaining hours 1-7 and 24 (0:00 – 7:00, 23:00 – 24:00).




Figure 9.1 Weekday averages of APX prices

Average APX prices on individual weekdays in the period 2 January 2001 till 14 June 2002 for baseload, peakload and off-peak hours.

The division in baseload, peakload and off-peak prices reflects part of the seasonality during a day, but prices also exhibit considerable seasonality during a week (see Figure 9.1). In general, prices (and electricity consumption) are lower during the weekend, especially on Sundays. At first sight there seems to be a downward trend from Monday to Friday, but this is probably due to some outliers⁵⁶. If we deseasonalize the natural logarithm of spot prices by reducing them with their weekday average, total variance is reduced on average with 20%: weekday influences explain about one fifth of total variance. In terms of our mathematical equation (2) $f(t)$ includes a dummy for Saturdays, and a dummy for Sundays (including public Holidays). In the Dutch spot prices we found only very weak evidence of seasonality over the year, so no specification is included for it. The seasonal component $f(t)$ is estimated jointly with the stochastic model parameters.

Prices in the Dutch market have witnessed already some serious spikes (see Figure 9.2). For example, the summer of 2001, which was expected to be a quiet period, contained some unexpected price movements. Those spikes even triggered an official investigation, but no irregularities were found: it was reported to be a simple coincidence of generator outages. Another series of high prices were observed close to the end of the year 2001. Since January 2002 prices were relatively stable till the second week of June, but June ended with some high prices.

⁵⁶ Median prices for each weekday are very stable from Monday to Friday, so the differences in average prices are mainly a result of a few outliers (or spikes).




Figure 9.2 APX baseload prices 2 January 2001 – 30 June 2002

APX prices are no exception to the phenomenon that prices in electricity markets are different from those in most other financial markets. This is clearly reflected in the summary statistics for the baseload, peakload and off-peak returns (see Table 9.1). All three series are characterized by a relatively high daily standard deviation of between 35 and 42%. For comparison: most individual stocks have daily standard deviations of 1-2% and only exceed 5% for the most risky stocks. The daily fluctuations in APX-prices can be enormous, reaching levels of over 200%. Especially the off-peak hours, with sometimes very low price levels, exhibit tremendous outliers. The extraordinary high kurtosis level of off-peak returns indicates that the fourth moment is probably not even defined. Returns are left-skewed, which may be surprising at first sight, since *prices* are clearly right-skewed due to the spikes. This indicates that prices do not only spike upwards very fast, but come down even faster, as the minimum and maximum returns indicate as well. This supports our choice to model spikes as a truly separate regime.

	Baseload	Peakload	Off-peak
Standard deviation	35.9%	39.5%	41.2%
Minimum	-204.4%	-186.2%	-452.3%
Maximum	161.3%	179.0%	396.3%
Skewness	-0.17	-0.07	-0.78
Kurtosis	4.60	3.80	45.19

Table 9.1 *Summary statistics APX returns*

This table presents summary statistics for the daily logreturns of APX day-ahead baseload, peak and off-peak indices in the period 2 January 2001 till 30 June 2002. Weekday influences (see Figure 9.1) were first removed from the price series before calculating the statistics.

9.2.2 Results

We use the sample with one and a half year of daily observations and three time-series to estimate model parameters. Even though the loglikelihoods of the regime-switching models are relatively complex, convergence was never a problem and independent of starting values. Results (Table 9.2) indicate that the regime models improve the fit considerably compared to the mean-reverting specification. Apparently, spikes that destroy the loglikelihood in the mean-reverting model are ‘transferred’ to the different spike and jump regimes where they fit considerably better.

Our model picks up on average 50% more spikes than the Huisman-Mahieu⁵⁷ model. This is most likely explained by the fact that their model requires that an up-jump is immediately followed by a down-jump, and is thus more restrictive on jumps. The speed of mean-reversion in both regime models is below that of the mean-reverting model; it is also lower for our model than for the Huisman-Mahieu model. Apparently, if strikes are not (or not so often) being detected, than a strong mean-reversion is required to pull prices back to normal levels. So the omission of spikes in the model specification leads to a misspecified mean-reverting process. Moreover,

⁵⁷ In the Huisman-Mahieu (2001) model we count the frequency of spikes as the sum of the up- and down-spikes.

since the regime models remove spikes from the stable process, their volatilities are considerably lower compared to the pure mean-reverting model. The mean-reverting volatility in our model is by far the lowest, because it transfers most erratic prices to the spike regime. The regime models indicate that the long-run average target levels for the baseload and peak spot prices are 5 and 8 €/MWh lower compared to the pure mean-reverting model.

Our regime specification picks up spikes well: expected spikes are positive for baseload and peakload, negative for off-peak, and have a much higher volatility than the stable mean-reverting process. When prices were mean-reverting on the previous day, a spike can be expected with a probability between 8 and 14%. Overall, between 15 and 27% of the prices are spikes, since spikes last on average around 2 days. For the pricing of far-out-of-the-money options, it might be considered that our regime model picks up small spikes too easily, and thus underestimates the magnitude of ‘real’ spikes. This is a common problem as well for stochastic jump models, and partly related to parameter calibration with maximum likelihood. It is possible to mitigate this problem with simple weights in the likelihood function⁵⁸.

The spikes in both regime models have an expected magnitude $E[S_t | \lambda(t) = S]$ of 53-57 €/MWh, and exhibit large possible swings. There is no doubt that the spikes deviate largely from the stable price levels and the data show it is crucial to separate them properly. For example, in the mean-reverting model a baseload price of over 100 €/MWh is nearly impossible (around 0.01% probability). In our regime-switch model its probability of 1.9% is much closer to the observed frequency of 2.6%.

For risk management purposes and derivative pricing, it is important that a model not only allows for spikes, but also for multiple consecutive spikes, a feature that is not contained in the Huisman-Mahieu model. In the estimation process the model assigns to each price a probability of being a mean-reverting price or a spike. The clustering of these ex post probabilities are informative to analyze, as we did in Figure 9.3 for baseload prices in the period June-August 2001. This was a turbulent period with several high spot prices, peaking even above 250 €/MWh on 3 July 2001. The graphs show several clusters of spikes. Examples are the periods 25-27 June, 2-7 July and 26-28 August 2001. Not surprisingly, these high spot prices are in general

⁵⁸ We tested that a very small adaptation to the loglikelihood function, that disfavors spikes, can reduce the frequency of spikes by over 50%, while at the same time keeping the loglikelihood within a few basis points from its maximum. We disfavoured spikes, by increasing the volatility in the $\ln(\sigma_t)$ term by 25% in the loglikelihood specification (8).

labeled a spike with a higher probability than low spot prices. For example, if a trader had sold a daily callable option for the first week of July, a considerable loss would have been his fate, and the outcome would have been overlooked by a risk management system that ignores the clustering of spikes over time.

Table 9.2 *Estimation results (next page)*

This table presents the parameter estimates and loglikelihoods of three different time series models for APX day-ahead baseload, peakload and off-peak hours. Estimates were obtained by maximum likelihood using data from 2 January 2001 till 30 June 2002. See the text for an explanation of all parameters and symbols.

		Baseload	Peakload	Offpeak
Mean-Reverting Model	α	0.414 (0.035)	0.421 (0.035)	0.576 (0.040)
	μ	3.414 (0.035)	3.609 (0.037)	2.795 (0.025)
	σ	0.323 (0.010)	0.353 (0.011)	0.307 (0.009)
	$E[S_t]$	32.018	39.292	17.154
	Sunday	-0.569 (0.035)	-0.613 (0.038)	-0.388 (0.035)
	Saturday	-0.231 (0.036)	-0.274 (0.039)	-0.058 (0.036)
	loglikelihood	-0.288	-0.377	-0.237
	Huisman and Mahieu (2001) Regime Model	α	0.404 (0.045)	0.399 (0.029)
μ_M		3.332 (0.030)	3.496 (0.028)	2.851 (0.028)
σ_M		0.207 (0.011)	0.209 (0.011)	0.155 (0.007)
$E[S_t \lambda(t) = M]$		28.589	33.727	17.509
μ_S		0.590 (0.202)	0.583 (0.089)	-0.668 (0.109)
σ_S		0.559 (0.136)	0.570 (0.053)	0.667 (0.069)
$E[S_t \lambda(t) = S]$		59.046	69.565	11.074
Sunday		-0.501 (0.025)	-0.528 (0.027)	-0.284 (0.019)
Saturday		-0.230 (0.026)	-0.266 (0.026)	-0.093 (0.018)
π_{MS}		0.066 (0.018)	0.092 (0.021)	0.071 (0.014)
% spikes		11.6%	15.6%	12.4%
loglikelihood		-0.109	-0.191	0.086
Mean-Reverting Regime Model with Spikes	α	0.356 (0.056)	0.243 (0.027)	0.170 (0.031)
	μ_M	3.289 (0.028)	3.433 (0.029)	2.858 (0.035)
	σ_M	0.157 (0.019)	0.123 (0.008)	0.116 (0.008)
	$E[S_t \lambda(t) = M]$	27.157	31.209	17.542
	μ_S	3.829 (0.131)	3.841 (0.058)	2.392 (0.077)
	σ_S	0.674 (0.065)	0.539 (0.039)	0.551 (0.048)
	$E[S_t \lambda(t) = S]$	57.760	53.826	12.728
	Sunday	-0.468 (0.027)	-0.474 (0.019)	-0.280 (0.021)
	Saturday	-0.217 (0.021)	-0.252 (0.019)	-0.100 (0.015)
	π_{MS}	0.084 (0.028)	0.139 (0.025)	0.110 (0.022)
	π_{SM}	0.473 (0.096)	0.385 (0.076)	0.599 (0.087)
	% spikes	15.1%	26.6%	15.6%
	Loglikelihood	-0.075	-0.137	0.166

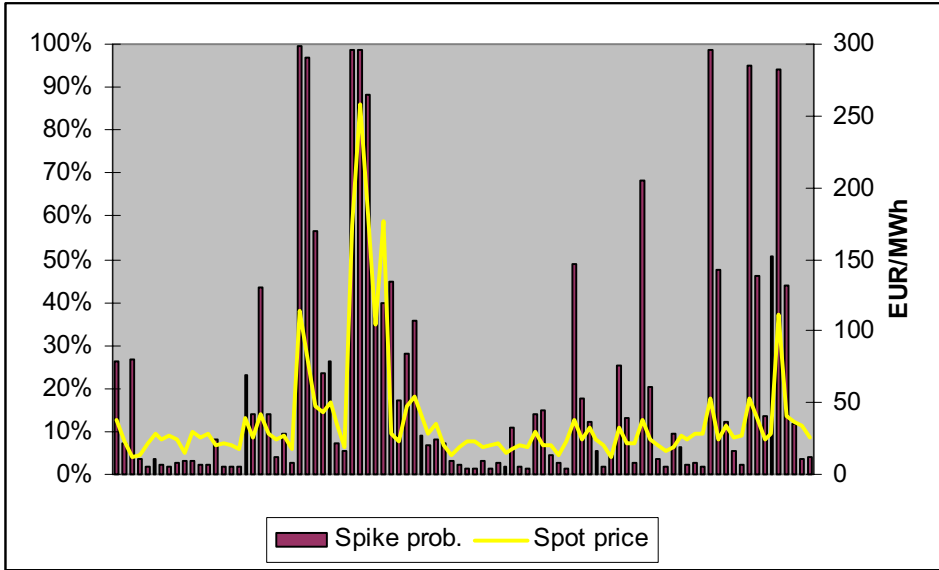


Figure 9.3 *Ex post spike probability*

Baseload APX prices and their ex post probability of being a spike in the period 1 June-31 August 2001.

9.3 Option valuation

In this section we present a methodology for the pricing of European-style options on spot prices using the proposed two-regime spot price model. We present closed-form formulas for standard calls and puts; pricing of caps, floors and swaptions is then straightforward. Closed-form formulas are important for various reasons. First, closed-form formulas may be more insightful than simulation-based calculations. For example, we will obtain two option value components: one related to the mean-reverting process, one to the spikes. Second, electricity traders often need to get quick answers in their day-to-day activities on the relative pricing of different options in the market. For them speed is often so important that it is necessary to use closed-form

formulas instead of simulation-based methods. Finally, closed-form formulas are very useful if options are being valued in a risk management application. Risk management statistics, such as Value-at-Risk can be computed much faster if no simulations are needed for the valuation of individual products in the portfolio.

Our results apply to European-style options on the spot price, which excludes a range of options that are traded in the marketplace. In electricity markets we also observe for example options on (average-price) forwards⁵⁹, and early exercise is sometimes allowed (American-style options). In the section with concluding remarks, we briefly describe how these kinds of options can be valued, within or without our regime framework.

9.3.1 Option valuation in a mean-reverting framework

For a better understanding of our approach, and to make meaningful comparisons between the two models, it is worthwhile to first review option valuation in a mean-reverting framework. Valuation of European-style options in a continuous time mean-reverting framework is for example described in Clewlow and Strickland (1999). A problem with continuous-time models is however that they cannot be perfectly estimated with discrete time data. Moreover, since a regime-switch model is inherently discrete (due to the regime switches), we have to make some adjustments to the aforementioned model.

In continuous time the mean-reverting model is formulated as follows (with dz_t a random draw from the standard Normal distribution):

$$\frac{dx_t}{x_t} = \alpha(\mu - x_{t-1})dt + \sigma dz_t \quad (14)$$

⁵⁹ In most markets a forward contract for the period of July-02 for example, entails the delivery of a constant electricity flow during the whole month of July at a fixed price. The value of the forward does therefore not only depend on the price on some particular day in July, but on the *average* price during the whole of July.

For the variance of x_τ (the variance of the log spot price τ periods from now) there exists a very elegant formula. Its properties are informative to analyze and similar to those of the variance in the discrete time model:

$$\text{Var}(x_\tau) = \frac{\sigma^2}{2\alpha} (1 - \exp(-2\alpha\tau)) \quad (15)$$

Without mean-reversion, the variance of a process converges to a level equal to the instantaneous variance multiplied by the time to maturity. With mean-reversion, the future distribution remains within stricter bounds: the higher the level of mean-reversion, the narrower is the future distribution. On forward prices this mean-reversion has the effect that the volatility of forward prices decreases exponentially from the current spot price volatility towards (almost) zero for longer maturities.

In Figure 9.4 the differences in volatility behavior are clarified between a standard Brownian motion and a mean-reverting time-series. In a world of Brownian motion the instantaneous volatility of a forward contract is independent of maturity⁶⁰. This culminates in a forward distribution at maturity with a standard deviation that is equal to the instantaneous volatility times the square root of maturity. In a mean-reverting model on the other hand, this standard deviation equals the square root of Equation (15) for a forward with maturity τ . Since the last term in (15) converges to unity for longer maturities, this standard deviation converges to a constant.

⁶⁰ See for example stock (index) futures whose volatility is nearly constant across maturity. This can simultaneously be explained by arbitrage and the non-existence of mean-reversion.

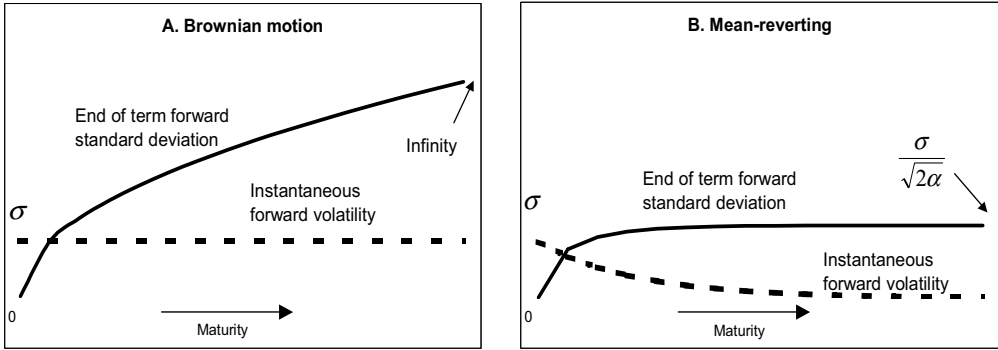


Figure 9.4

Forward volatilities of a Brownian motion (A) and a mean-reverting (B) continuous time process. The instantaneous volatility is the volatility of the forward at time 0. The end of term volatility is the standard deviation of the forward return from time 0 till maturity.

Returning to the discrete time world, no elegant formulas as Equation (15) for the variance are available. Instead, the variance of the process must be determined recursively from the variance one period back. The mean-reverting model assumes that the one period ahead variance equals σ^2 . By taking the variance of Equation (4) we obtain the variances of x_τ for maturities $\tau = 1, 2, \dots$:

$$\begin{aligned} Var(x_1) &= \sigma^2 \\ Var(x_2) &= (1 - \alpha)^2 \sigma^2 + \sigma^2 \end{aligned} \tag{16a}$$

....

$$\begin{aligned} w_\tau = Var(x_\tau) &= (1 - \alpha)^2 Var(x_{\tau-1}) + \sigma^2 \\ w_\infty &= \frac{\sigma^2}{1 - (1 - \alpha)^2} \end{aligned} \tag{16b}$$

Similar to the variance in the continuous time model, this variance converges to a constant as well (Equation 16b). For small mean-reversion parameter α the difference between the two approaches zero, but for increasing levels of mean-reversion, it pays

off not to treat a discrete time model as a continuous time model. For example, with mean-reversion parameter α equal to 1/2, the stationary variances are a non-negligible factor 1/3 apart.

If spot prices are mean-reverting (so without spikes yet), then all ingredients to value options on forward contracts are now available (Clewlow and Strickland, 1999): an end-of-term forward distribution being lognormal, market forward prices and the variance as defined in Equation (16). Application of the Black's (1976) formula for the valuation of a European call option with maturity τ and strike price K thus yields ($N(\cdot)$ is the standard normal cumulative distribution function):

$$d = \frac{\ln\left(\frac{F_\tau}{K}\right) + \frac{1}{2} w_\tau}{\sqrt{w_\tau}} \quad (17a)$$

$$Call(\tau, K) = \exp(-r\tau) \cdot \left\{ F_\tau N(d) - KN\left(d - \sqrt{w_\tau}\right) \right\} \quad (17b)$$

9.3.2 Option valuation in the regime switch model

The idea behind the option valuation in the proposed regime switch model, is to split up the option price in a mean-reverting component and a spike component. We use the feature that the prices in both regimes are independent and lognormally distributed, even though the volatilities of both regimes may be wide apart. The volatility of the mean-reverting price distribution changes with maturity, as explained earlier, but the spike distribution is independent of maturity. Since the prices in the two distributions are independent an option value can be calculated for each regime, conditional on the price process being in that regime. The weight that each component receives, is determined by the probability of a spike. The reason that the two option components can be added up, is that the spikes are independent from the mean-reverting prices. For example, a European-style call option on the spot P_t , with maturity τ and strike K has a 'fair' value of (ignoring time-value):

$$\begin{aligned}
Call &= E[\max(0, P_\tau - K)] \\
&= \pi_{M,\tau} \cdot E[\max(0, P_\tau - K) | \lambda = M] + \pi_{S,\tau} \cdot E[\max(0, P_\tau - K) | \lambda = S] \quad (18) \\
&= \pi_{M,\tau} \cdot Call_M + \pi_{S,\tau} \cdot Call_S
\end{aligned}$$

where $\pi_{M,\tau}$ and $\pi_{S,\tau}$ are the prior probability of a mean-reverting regime and spike regime at maturity τ respectively and λ indicates the type of regime. $Call_M$ and $Call_S$ are the values of a call option if the process would be in the mean-reverting and spike regime. It is important to understand that the value of a call option is a weighted average of these two regime-dependent call options, because τ periods from now, the spot price is not a weighted *combination* of mean-reverting price and spike price, but *either* a mean-reverting price *or* a spike price⁶¹. How to calculate the values of the two regime-dependent call values is defined below.

Variances in each regime

For the valuation of the mean-reverting component ($Call_M$), we apply Black's (1976) formula as explained in the previous section (Equation 16 and 17). For the valuation of the spike-component ($Call_S$) we apply Black's (1976) result again, because the spikes are lognormally distributed as well. The inputs to both formulas are the forward prices and the variances of the price processes in the two regimes. The variance of the log spot price in the spike process is independent of maturity and equals σ_S^2 . For the variance of the mean-reverting component we rely on the recursive formula in Equation (16a).

$$w_{S,\tau} = Var(x_{S,\tau}) = \sigma_S^2 \quad (19a)$$

$$w_{M,\tau} = Var(x_{M,\tau}) = (1 - \alpha_M)^2 w_{M,\tau-1} + \sigma_M^2 \quad (19b)$$

⁶¹ The expected spot price however is a weighted combination.

Forward price levels in each regime

In order to determine the forward price in the mean-reverting and spike regime, the expected spot prices in both regimes could be used. However, it is preferable to take account of market forward prices, because we don't want our derivative values to deviate from market prices because we have a different view on forward prices, but because we have a different view on the spot price process, such as its volatility and level of mean reversion. Therefore, the spot price level needs to be aligned with market expectations. Moreover, this avoids the tedious modeling and estimation of all seasonal influences and risk premia. For example, in our spot price data we could not find significant seasonal variations over the year, but the current forward curve indicates that traders believe prices in the winter to be higher than in the summer.

While employing market forward prices, we have the choice to adjust either the expected spike level to market forward price levels or to adjust the mean-reverting level or to adjust the relative probabilities of the regimes (or both). Each of these changes is defensible. We choose to adjust the expected spike level, because we believe it's especially the risk of spikes that justifies risk premia in forward markets. Such a risk premium may cause a possible mismatch between market forward prices and expected spot price levels in our model.

Our procedure works as follows. Suppose we observe market forward prices F_τ with maturities $\tau = 1, \dots, N$. Then we need to find the appropriate spike and mean-reverting forward prices $F_{S,\tau}$ and $F_{M,\tau}$ such that the probability-weighted sums equal the market forward prices:

$$F_\tau = \pi_{S,\tau} F_{S,\tau} + \pi_{M,\tau} F_{M,\tau} \quad (20)$$

where $\pi_{M,\tau}$ and $\pi_{S,\tau}$ are the posterior probabilities of a mean-reverting regime and spike regime at maturity τ respectively. These probabilities of a future regime depend on the regime today and the regime switching probabilities. They can be calculated by τ times premultiplying today's posterior regime probabilities with the transition matrix Π , defined in Equation (5). We then determine recursively the mean-reverting forward price level, based on the result that it equals the expected spot price level in the mean-

reverting regime (see Equation 21a below). The remaining part of the market forward price belongs to the spike regime (Equation 21b).

$$\begin{aligned}
 \ln F_{M,\tau} &= \ln E[P_\tau | \lambda = M] \\
 &= \ln E[\exp(x_{M,\tau} + f(\tau))] \\
 &= E[x_{M,\tau}] + \frac{1}{2} \text{Var}[x_{M,\tau}] + f(\tau) \\
 &= \alpha_M \mu_M + (1 - \alpha_M) E[x_{\tau-1}] + \frac{1}{2} \text{Var}[x_{M,\tau}] + f(\tau) \\
 &= \dots\dots\dots + \frac{1}{2} \text{Var}[x_{M,\tau}] + f(\tau)
 \end{aligned}
 \tag{21a}$$

$$F_{S,\tau} = \frac{F_\tau - \pi_{M,\tau} F_{S,\tau}}{\pi_{S,\tau}}
 \tag{21b}$$

Option values in each regime

With their means and variances specified, the lognormal distributions of the spot price processes are now completely defined. The option values conditional on each regime, $Call_M$ and $Call_S$ are obtained with the Black (1976) formula, as defined in equation (17). We thus obtain:

$$Call(K, \tau) = \pi_{M,\tau} \cdot Call_M(K, \tau) + \pi_{S,\tau} \cdot Call_S(K, \tau)
 \tag{22}$$

where

$$Call_M(\tau, K) = \exp(-r\tau) \cdot \{F_{M,\tau}N(d_M) - KN(d_M - \sqrt{w_{M,\tau}})\}$$

$$Call_M(\tau, K) = \exp(-r\tau) \cdot \{F_{S,\tau}N(d_S) - KN(d_S - \sqrt{w_{S,\tau}})\}$$

$$d_M = \frac{\ln\left(\frac{F_{M,\tau}}{K}\right) + \frac{1}{2}w_{M,\tau}}{\sqrt{w_{M,\tau}}}$$

$$d_S = \frac{\ln\left(\frac{F_{S,\tau}}{K}\right) + \frac{1}{2}w_{S,\tau}}{\sqrt{w_{S,\tau}}}$$

This completes the derivation of a European-style option value on a spot electricity price, where the exercise price is K and the option matures at maturity τ .

9.3.3 Option valuation example

As an example we apply the above-described procedure to the forward curve in the Dutch market on 30 June 2002. Forward contracts in electricity markets apply to delivery of the commodity *during* a certain period. So, the forward prices depend on the expected *average* spot prices in those periods. In our model however, we need the price of a forward maturing on one single day. We calculate the value of options on the spot that mature in the middle of each delivery period and calculate the number of days (τ) till those dates (see Table 9.3). With the moderate preceding weekend prices, the model produces a spike probability of 9.0% on the first day of July for baseload prices and 15.5% for peakload prices. These prior probabilities are obtained by pre-multiplying them iteratively by the transition matrix Π , and converge relatively fast to a stable 15.1% for baseload and 26.6% for peakload.

Next, we derive the forward prices and variances for each regime. The model (Equation 21a) determines the mean-reverting forward price. These mean-reverting forward prices are some Euros below the market forward prices (Table 9.3). The difference between the two is due to the expectation of occasional spikes (Equation

21b), which lie in the range of 50-90 €/MWh for baseload and 75-120 €/MWh for peakload. The variance of the spike forward price is stable across maturity, but the variance of the mean-reverting forward converges to a level of 4.2% for baseload and 3.5% for peakload.⁶² This is considerably lower than the spike variances of 45.5% and 29%. Therefore, combined with their considerably higher expected level, it is not surprising that the spikes form an important ingredient of the option values.

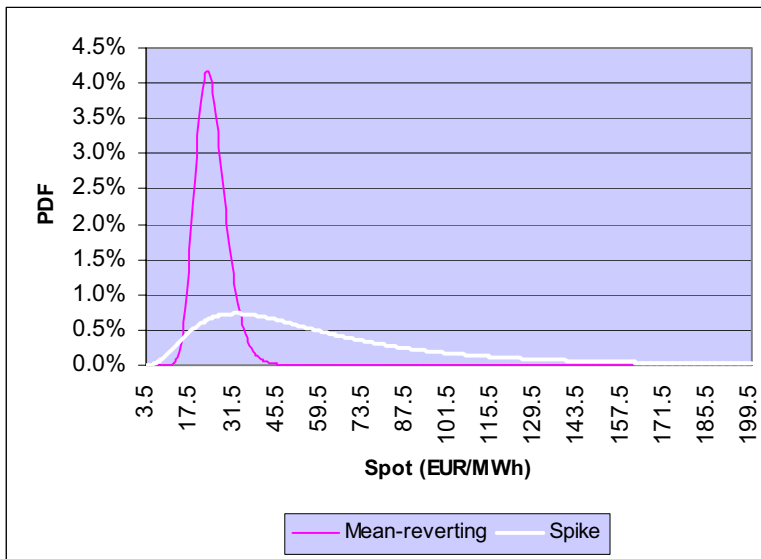


Figure 9.5 *Mixture of lognormals*

Probability density functions for the spot price regimes on 15 July 2002, based on the model estimated on 30 June 2002.

⁶² The slightly higher mean-reverting variance for baseload than peakload prices may be explained by the relatively high frequency of spikes obtained from the peakload estimates.

Contract	τ	Price	Pr[Spike]	$F_{M,\tau}$	$F_{S,\tau}$	$w_{M,\tau}$	$w_{S,\tau}$
Panel A: Baseload contracts							
Day ahead	1	27.00	9.0%	22.27	74.57	2.5%	45.5%
Week ahead	7	32.50	15.0%	24.17	79.59	4.2%	45.5%
July	15	30.25	15.1%	24.32	63.66	4.2%	45.5%
August	46	28.43	15.1%	24.32	51.56	4.2%	45.5%
Q4-02	138	30.13	15.1%	24.32	62.84	4.2%	45.5%
Q1-03	230	31.55	15.1%	24.32	72.26	4.2%	45.5%
Q2-03	319	34.25	15.1%	24.32	90.17	4.2%	45.5%
Panel B: Peakload contracts							
Day ahead	1	38.00	15.5%	25.79	91.58	1.5%	29.0%
Week ahead	7	41.75	26.4%	27.51	81.41	3.5%	29.0%
July	15	40.75	26.6%	27.88	76.36	3.5%	29.0%
August	46	41.25	26.6%	27.92	78.12	3.5%	29.0%
Q4-02	138	43.00	26.6%	27.92	84.71	3.5%	29.0%
Q1-03	230	46.25	26.6%	27.92	96.95	3.5%	29.0%
Q2-03	319	52.13	26.6%	27.92	119.09	3.5%	29.0%

Table 9.3 *Forward curve construction*

This table describes the process of moving from market forward prices to a spot price model that is in line with the market. The first column contains the contract type, the second column the time-to-maturity in days. The third column lists the Dutch forward prices in €/MWh on 1 July 2002 for baseload and peakload electricity respectively. Source: Platts Dutch Power Assessments delivered through Moneyline/Telerate. The other columns combine the parameter estimates with the market prices to show the probability of a spike (Pr[Spike]), and the forward prices and variances of each regime.

In Figure 9.5 we plot the distributions of each regime for expected 15-day ahead baseload prices, based on the parameter estimates in Table 9.2. The actual distribution is a weighted average of the two individual regime-dependent probability density functions. In the example this means approximately 85% weight for the mean-reverting and 15% for the spike regime. Although a 15-day maturity is not particularly long, we observe in Figure 9.5 that the mean-reverting regime prices have a very narrow distribution, which is due to a low daily standard deviation of 20.5%. The expected mean-reverting spot price is with 99% probability in the range of 14 to 40 €/MWh. The spike regime prices in contrast are much wider distributed with a standard deviation of 67.5%. Here a 99% confidence interval covers a range as wide as 9-288 €/MWh. Prices reaching very high levels are thus far more likely in the regime-switch model (where they occur in the spike regime) than in the mean-reverting model (without spikes).

In the mean-reverting model the standard deviation of the expected spot price is somewhere in between the mean-reverting regime and the spike regime at a level of 39.8%. This implies that spot prices 15 days in the future will leave a bandwidth of 20-46 €/MWh only once in 100 years. That's why the option values that result from our regime-switch spot price model largely deviate from options in a mean-reverting framework. In Table 9.4 we make a comparison. Option values were calculated of call options that mature in 1, 7, 15 and 46 days, with strike prices of 20, 30, 40 and 50 €/MWh. We take the parameter estimates and forward values of 30 June 2002 (Table 9.2 and 9.3) and assume an interest rate of 4%.

Mean-Reverting Model					Regime-Switch Model			
Strike (€/MWh)								
Maturity	20	30	40	50	20	30	40	50
Panel A: Baseload options								
1	7.71	2.34	0.57	0.13	7.46	4.18	3.43	2.84
7	12.50	4.00	0.61	0.06	12.90	7.96	6.40	5.35
15	10.27	2.59	0.28	0.02	10.66	5.73	4.30	3.43
46	8.46	1.68	0.14	0.01	8.89	4.09	2.83	2.14
Panel B: Peakload options								
1	18.13	9.77	4.52	1.90	18.01	11.72	10.08	8.68
7	21.73	11.87	4.14	0.86	21.79	14.47	11.31	9.18
15	20.71	10.87	3.41	0.59	20.77	13.32	10.11	8.05
46	21.12	11.30	3.68	0.67	21.18	13.73	10.50	8.41

Table 9.4 *Call option values*

This table shows values of European-style options on power spot prices according to the mean-reverting model (panel A) and the regime-switch model (panel B). Parameter estimates are taken from Table 9.2 and forward prices from Table 9.3. Option values are calculated for four different maturities and strikes of 20, 30, 40 and 50 €/MWh.

For low strike prices, there is very little difference between the mean-reverting and the regime-switch option values. Their main value driver for low strikes prices is the current forward price, which explains the differences between maturities. Increasing the strike level progressively from 20 €/MWh to 50 €/MWh option values in the mean-reverting model quickly decline towards zero. As we saw earlier, there is not much weight in the right tail of the mean-reverting distribution, resulting in hardly any value for deep out-of-the-money options. In the regime-switch model however, the spikes take account of the right tail, which explains that option values are substantial even for far out-of-the-money options. This clearly leads to more realistic option values that take into account that spot prices can be very erratic.

The difference between the two models is best understood if we consider options with a strike price of 50 €/MWh. The mean-reverting model indicates those options are close to worthless, although we know that it is certainly not impossible

that spot prices will reach levels above 50 €/MWh on individual days. As an illustration, in our sample baseload prices exceeded this level on more than 8% of the days and peakload prices on more than 13% of the days. So even options with high exercise prices have substantial value, which is entirely ignored by the mean-reverting model.

For example, the costs of a maximum price (cap) would severely be underestimated with the mean-reverting model. Caps are equal to a series of call options and frequently embedded in retail electricity contracts, where they form a bridge between fixed and floating price contracts. Let's consider a contract where the end-user pays the daily baseload APX-price on each day in July, but with a cap of 50 €/MWh. If we take the possibility of spikes into account, such a cap would cost approximately 3.43 €/MWh (based on an average maturity of 15 days), whereas a supplier would give it away for free if the wrong model were being used.

9.4 Concluding remarks

In this chapter we presented a model to value options on electricity spot prices. It takes into account the two main features of electricity prices: strong mean-reversion and occasional 'spikes'. Closed-form formulas for European-style options were obtained by disentangling the mean-reverting spot prices from the spikes, such that option values can be broken down in two components that were each valued with Black's (1976) formula for options on forwards and futures. We showed that it is crucial to include spikes in any option price formula, since they represent substantial value, especially for far out-of-the-money options.

Our results apply to European-style options on the spot price, which excludes a range of other tradable options. In electricity markets we also observe for example options on individual hours, options on forwards, and early exercise is sometimes allowed (American-style options). We believe a regime model could work well for individual hours, since these are characterized by periods of spikes and stable periods of mean-reverting prices as well. A difficulty here is that individual hours are even more volatile and spiky than daily averages, exhibit strong seasonality, and that prices on the same day strongly interact.

Forwards whose value depend on the average price during a certain period on the other hand are hardly affected by the presence of spikes, as long as the averaging

period is long enough. The pricing of options on those forwards can therefore best be done by modeling the forwards directly, instead of aggregating spot prices into average price forwards. Moreover, the main uncertainty for options on forwards emanates from variations in the long-run average price. Our model is a one-factor model, which means that the long-run average price is not stochastic. This is not so relevant for the valuation of options on the spot, but it would be incorrect to use the same one-factor model for long-term options on forwards.

Early exercise is not really an issue in valuing options on spot prices. Naturally, the holder of the option will wait till the last moment before deciding to exercise or not. Therefore, in practice American-style options on the spot will behave like European-style options.

Our application of Black's formula implies that the risks of the mean-reverting prices and the spikes can be hedged, which is not completely realistic. With increasing liquidity in most electricity forward markets, it might be possible to hedge some part of the uncertainty in option prices, but the largest uncertainty of options on the spot result from the spikes, and there are no financial⁶³ strategies to hedge spikes properly. Therefore, the derived option values can best be regarded as fair prices if uncertainty is ignored⁶⁴. However, market participants might be willing to price options somewhat higher, because spikes make especially selling options risky. The model makes such an adjustment relatively easy, since it yields an explicit value for the spike component of the option value, which may be adjusted to include a risk premium.

The separation of the spikes from the mean-reverting prices ensures that only a limited number of parameters needs to be estimated. This is important in electricity markets where we have only a relatively short history of reliable prices, and markets are in constant change. As markets become more mature, it may be worthwhile to include more electricity price characteristics, such as time-varying volatility and time-varying spike intensities. That will be the subject of future research.

⁶³ A way to hedge spikes physically is by keeping some reserve capacity to use when prices are unexpectedly high.

⁶⁴ While using market forward prices, we incorporated the risk premium in the forwards. Options might however justify an even larger risk premium.

10 Conclusion of the second part

In the last three chapters we presented three different option pricing methodologies that all deviate from the standard assumption of normally distributed returns. We end this second part of the thesis with a brief comparison and an analysis of their applicability in real-world markets.

Since the stock market crash in 1987 option prices seem to deviate from the assumption of constant volatility across strike and maturity. A large number of researchers have since investigated explanations for the observed skews, smiles and term structure effects. In the chapter 7 and 8, we explored two econometric methodologies that assume that these deviations can be explained by different expected distributions or price processes than the standard normal (or Brownian motion). Other plausible explanations were not directly dealt with in the two chapters. For example, transaction costs, different risk premia across strike or maturity, difficulties to hedge options properly, and general inefficiencies in market prices may just as well explain implied volatility patterns. These alternative explanations are for example studied in Jackwerth (2000), and Ait-Sahalia, Wang and Yared (2001). Indirectly however, we did investigate the role of these alternative explanations and even found some support for them. In Chapter 7 for example it appeared that methods which model the volatility curve obtained a better fit to market option prices than methods that model the underlying risk-neutral distribution, such as the skewed Student-t method we proposed. This result may lead to the conclusion that the distribution-methods analyzed are not flexible enough, but we believe that the focus of option traders on volatility numbers is a more creditworthy explanation. Volatility numbers are easier to understand and better comparable across different options, and therefore an important statistic in day-to-day trading. We believe a small ‘inefficiency’ in market prices is the result of this, although exploiting it may be hard or impossible with bid-ask differences and other transaction costs in place.

In Chapter 8 we investigated the information content of option prices relative to time-series data in forecasting short-term volatility. Although a more direct test on the efficiency of option prices, we did not label it an efficiency test, because we incorporated time-series information (risk parameter and long term volatility) into the

implied estimates as well. We did so because both options and time-series may contain information about future prices. Options are explicitly forward looking and may respond quickly to news regarding the future price process. Option prices are on the other hand not solely determined by the expected future distribution, but also by risk preferences and other ‘inefficiencies’. The implied GARCH method we proposed aims at using both information sources to its best, but other approaches that combine option and time-series information may also perform well. Especially promising in this area are approaches that combine both high-frequency data with implied option data. From a practitioner’s viewpoint however, there are still many hurdles to take. Huge data sets, careful data management and complicated estimation procedures of only the most liquid instruments will yield satisfactory outcomes. Application in day-to-day business and beyond a few currency pairs, stock indices and commodities seems a long way to go.

In electricity markets it is equally tempting, but not always practical due to data limitations, to employ the most advanced techniques in option pricing. In Chapter 9 we proposed a relatively parsimonious model that describes electricity spot prices. Its regime switches build upon the foreign exchange literature, where this type of model has originally been proposed. The model fits electricity markets well, because the lack of storage opportunities makes spot prices largely dependent on demand and supply conditions at that particular point in time, creating different regimes on different days. On most days price formation is a rather predictable stable mean-reverting process, but every now and then generation problems, network congestions or unexpectedly high demand cause prices to spike to very unpredictable levels. In a market that is still in a very early development phase, we hope our model adds to the correct management of uncertainty and pricing of real and financial assets. Correct, reliable and practical models may ultimately lead to greater confidence, lower risk aversion and a better functioning of markets in general.

11 Summary and concluding remarks

This thesis provides the results of five studies on financial derivatives. In this final chapter we briefly summarize the results of each study and formulate some general conclusions.

11.1 Summary first part

In the first part of this thesis we investigate the impact of options trading on the time-series properties of the underlying asset. More precisely, we test the hypothesis that a derivative asset improves the efficiency in the underlying asset such as a stock. The two studies presented largely support this hypothesis, and clarify the mechanisms that lead to this result. A plausible explanation is that the presence of a correlated asset permits the sharing of effective price discovery across markets. Market makers in the stock can set more accurate prices if they learn from transactions in the option.

In Chapter 3 we use a controlled trading environment, where students trade in a stock and a call option on the stock in markets with asymmetric information. This allows the observation of all information sets and all actions in a setting based on the Kyle (1985) framework, but beyond the reach of tractable modeling. Repeated trading rounds with different groups of students and different asset values make clear that an insider trades aggressively in both the option and the stock, with most trades directed to the asset that affords the most profitable trading opportunity. This leads to price discovery occurring in both markets, and hence important feedback effects: trades in the stock market imply quote revisions in the options market and vice versa. Because the comparison of a market with and a market without options trading leads to control problems (due to a different number of market participants) we decided to study the effect of options trading indirectly by analyzing the impact of the option's moneyness. We find a significant impact of moneyness on the time-series properties of the underlying: when the option is in-the-money the convergence to informationally efficient pricing is more rapid and the volatility of transaction prices is lower. The tendency of insiders to trade where the magnitude of the profitable trading opportunity

is greatest, provides a richer set of signals to dealers than when there is only a single asset in which the insider can trade profitably. We show therefore that not only does the presence of a correlated asset effectively split price discovery across markets; it also fundamentally changes the process by which conditional expectations are updated. We furthermore show that the less strategic the insider (due to risk-aversion, impatience, or noisy signals), the more powerful we expect this effect to be.

Our model in Chapter 4 provides two important extensions to the existing theoretical models. First, we extend a single-trade model to a dynamic multi-trade environment. Second, we analyze market quality under different levels of option leverage, the main distinguishing property of options. We start with a standard sequential trade model, and show that it is inherently dynamic. Because expectations are updated after every trade, we can simulate a sequence of trades and derive more precise criteria for market quality. Our model clarifies and separates two mechanisms following the introduction of an option. On the one hand, stock dealers learn from trades in the option market and set more accurate prices. On the other hand, the proportion of informed trading in the stock is altered depending on the option's effective leverage, possibly reducing some market quality statistics. Our dynamic model indicates that option trading always reduces price errors (difference between the intrinsic value and the traded price) in the underlying, because an option serves as an extra source from which information can be inferred. Uninformed traders benefit most from this reduction in price errors in a derivatives market that allows for relatively large (informed) trades, and in a market where the number of informed traders is relatively small. This corresponds to well-developed derivatives markets where options provide important leverage. In terms of price volatility, effective leverage has the opposite effect: trading in well-developed derivatives markets leads to higher volatility.

11.2 Summary second part

In the second part of this thesis we present three empirical studies on option pricing for assets with non-normal returns. In the first two studies we explore methods to infer information from market option prices. The last chapter is a more standard pricing study, but applied to a very non-standard and risky commodity, electricity.

Chapter 7 presents a methodology to derive the *risk-neutral* distribution from option prices in a flexible and accurate manner. The exact shape of the implied risk-

neutral distribution gives important information that can be used for pricing other options on the same underlying asset, for comparing options on different assets and for closely monitoring changes in the markets perception of the underlying price process. We apply a skewed version of the well-known Student-t distribution to capture the smiles and skews frequently observed in implied volatility curves. Its main strength is the direct parameterization of skewness and kurtosis. In an application to several years of FTSE-100 index options we compare the in-sample performance of the skewed-t method with the normal method (constant volatility), two implied volatility curve-fitting methods and a trinomial tree. Although they all clearly improve upon the normal method, the volatility curve-fitting method that regresses implied volatility on option delta outperforms the trinomial tree and skewed-t methods: average root mean squared errors are lower and this effect strengthens over time. We conclude that a curve-fitting method with the option's delta as explanatory variable is preferred to price European-style options outside the available trading range. However, even though their fit is inferior, the two methods that focus on modeling the distribution of asset returns do have strong appeals. The skewed-t method in particular has the appeal that its parameters relate directly to the moments of the distribution. This makes it possible to accurately monitor changes in market expectations about the underlying asset.

Chapter 8 presents a method to infer from option prices a forecast of the actual price process, instead of the distribution at a single point in time. We use the Duan (1995) option-pricing model to identify an option-implied EGARCH process and the corresponding 1-day ahead volatility forecast. For this implied process we estimate the stable long-term volatility and the risk premium from the time-series of daily index returns. We compare the implied forecast to the forecast of a (pure) time-series EGARCH model on FTSE-100 and DAX-30 index data. Our results yield different outcomes on the two markets. In the German market (DAX-30), the implied volatility forecast is leading the historical forecast, but no such relation is found in the UK market (FTSE-100). In the German market, the historical forecast explains around 23% of realized volatility (constructed from intraday data), whereas the same ability of the implied estimate is 1.5 times larger. In the UK market both predictors perform equally well and explain a large fraction of around 44%. In the UK market a weighted combination of the two predictors explains even 8% extra, but in the German market the implied forecast alone contains nearly all information. Finally, in the German market the implied forecast is a relatively unbiased predictor of realized volatility

including overnight returns, but not so in the UK. So even though the implied forecast in the UK explains a large fraction of actual volatility, it does no better than a time-series EGARCH model. In the German market on the other hand, the implied forecast is clearly preferred.

Chapter 9 presents a regime-switch model for the pricing of options on spot electricity prices. The spot price model incorporates the main features of electricity prices: seasonality, mean-reversion, high volatility and occasional spikes. A major contribution of the spot price model is that it allows for multiple consecutive spikes, which is important for risk management and derivative pricing purposes. Based on this spot price model we present closed-form formulas to price European-style options on spot electricity prices. Because the regime-switch model contains two independent regimes, option prices can be split up in two components: one for the stable mean-reverting process, one for the spikes, which can both be valued with the Black (1976) model. Application to the Dutch APX market shows the importance to include spikes properly in valuing options: option prices are considerably higher and closer in line with historical pay-offs with the regime-switch model than with a mean-reverting model. This effect is especially strong for the popular out-of-the-money call options.

11.3 Concluding remarks and future research

Both studies in the first part of this thesis reveal the complex relations and dependencies with trade in two correlated assets. The experimental study indicates that efficiency improves, but the theoretical study clarifies that part of this result may depend on the exact characteristics of a market. By abstracting from real human behaviour, we believe that the theoretical model somewhat underestimates the benefits of derivatives. In the model all market participants draw exactly the same conclusions from each trade. As a result, there are no differences in opinion among stock and option dealers. In the experiments, similar to real-world markets, those differences in opinion (and the insider response) are particularly informative, and speed up the price discovery process. For a better understanding of human behavior in a market with correlated assets, future research should therefore be devoted to the development of models that allow for more realistic human behavior. Alternatively, new experiments can be set up with more variations in control variables.

Regulators that control derivatives markets may use the results of the two studies and future research to better set the standards for derivative markets. For

example, regulators need to weigh the benefits of lower losses for uninformed traders against possibly increased market volatility. They furthermore need to decide on the effective leverage derivative markets may optimally provide. According to the two studies, the benefits of derivative trading to improved price convergence are however without doubt.

In the second part of the thesis we build upon the positive conclusion of the first part and analyze three different option-pricing methodologies for three different applications. Although the seminal work of Black, Scholes and Merton has greatly stimulated the development of derivative markets, we show that the assumptions in their models need to be adjusted for more realistic option pricing. With a recent modeling approach we show for example that options may contain valuable information beyond the information in time-series returns. In the very peculiar electricity markets we develop a model for spot electricity prices, derived from a popular foreign exchange model, that yields realistic option prices. Since option pricing techniques and data quality are constantly evolving, future research based on the second part of the thesis is relatively easy to formulate. In mature markets especially the incorporation of high-frequency data is a promising way to go, with the danger however that results can be applied to only a few highly liquid commodities. In electricity markets option pricing and risk management are still in their infancy and many new techniques need to be developed to price the different types of assets. With better models we aim to improve the quality of derivatives markets, and enhance confidence among participants that should ultimately lead to a better diversification of risks.

References

- Aït-Sahalia, Y. and A. Lo (1995), "Nonparametric estimation of state-price densities implicit in financial asset prices", *Working paper series 5351*, National Bureau of Economic Research
- Aït-Sahalia, Y., Y. Wang and F. Yared (2001), "Do option markets correctly price the probabilities of movement of the underlying asset?", *Journal of Econometrics*, 102, 9-51
- Akgiray, V. (1989), "Conditional heteroskedasticity in time series of stock returns: evidence and forecasts", *Journal of Business*, 62, 55-80
- Amin, K I. (1993), "Jump diffusion option valuation in discrete time", *Journal of Finance*, 48, 1833-1863
- Amin, K I., and Charles M.C. Lee (1997), "Option trading, price discovery, and earnings news dissemination", *Contemporary Accounting Research* 14, 153-192
- Anagnou, I., M. Bedendo, S. Hodges and R. Tompkins (2001), "The relationship between implied and realised probability density functions", *working paper*, University of Technology, Vienna
- Andersen, T.G. and T. Bollerslev (1998), "Answering the skeptics: yes standard volatility models do provide accurate forecasts", *International Economic Review*, 39, 885-905
- Andersen, T.G., T. Bollerslev, F.X. Diebold and H. Ebens (2001a), "The distribution of realized stock return volatility", *Journal of Financial Economics*, 61, 43-76
- Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2001b), "The distribution of exchange rate volatility", *Journal of the American Statistical Association*, 96, 42-55
- Back, K. (1993), "Asymmetric information and options", *Review of Financial Studies* 6, 435-472

Ball, S.B., M.H. Bazerman, and J.S. Carroll (1991), "An evaluation of learning in the bilateral winner's curse", *Organizational Behavior and Human Decision Processes* 48, 1-22

Bates, D.S. (1991), "The crash of '87: Was it expected? The evidence from option markets", *Journal of Finance*, 46, 1009-1044

Bates, D.S. (1996), "Testing option pricing models", in Maddala, G.S. and C.R. Rao (eds.), *Handbook of statistics 1.4: Statistical methods in finance*, Elsevier Science, North-Holland, 567-611

Bauwens, L. and M. Lubrano (2002), "Bayesian option pricing using asymmetric GARCH models", *Journal of Empirical Finance*, 9, 321-344

Bhanot, K. (2000), "Behavior of power prices: implications for the valuation and hedging of financial contracts", *Journal of Risk*, 2:3, 43-62

Biais, B., and P. Hillion (1994), "Insider and liquidity trading in stock and options markets", *Review of Financial Studies* 74, 743-780

Bjerksund, P., Rasmussen, H. and G. Stensland (2000), "Valuation and risk management in the Norwegian electricity market", *discussion paper 20/200*, Norwegian School of Economics and Business Administration

Black, F. and M. Scholes (1973), "The pricing of options and corporate liabilities", *Journal of Political Economy*, 81, 637-659

Black, F. (1975), "Fact and fantasy in the use of options", *Financial Analysts Journal*, 31, 36-72

Black, F. (1976), "The pricing of commodity contracts", *Journal of Financial Economics*, 3, 167-179

Blair, B.J., S.-H. Poon and S.J. Taylor (2001), "Forecasting S&P100 volatility: The incremental information content of implied volatilities and high frequency index returns", *Journal of Econometrics*, 105, 5-26

- Bliss, R.B. and E.I. Ronn (1989), "Arbitrage-based estimation of nonstationary shifts in the term structure of interest rates", *Journal of Finance*, 44, 591-610
- Bloomfield, R. and M. O'Hara (1998), "Does order preferencing matter?", *Journal of Financial Economics* 50, 3-37
- Bloomfield, R. and M. O'Hara (1999), "Market transparency: Who wins and who loses?", *Review of Financial Studies* 12, 5-35
- Bloomfield, R., and M. O'Hara (2000), "Can transparent markets survive?", *Journal of Financial Economics* 55, 425-459
- Bollerslev, T. (1987), "A conditionally heteroskedastic time series model for speculative prices and rates of return", *Review of Economics and Statistics*, 69, 3, 542-547
- Bollerslev, T., R.Y. Chou and K.F. Kroner (1992), "ARCH modeling in finance: A selective review of the theory and empirical evidence", *Journal of Econometrics*, 52, 5-59
- Brailsford, T.J. and R.W. Faff (1996), "An evaluation of volatility forecasting techniques", *Journal of Banking and Finance*, 20, 419-438
- Breedon, D. and R. Litzenberger (1978), "Prices of state-contingent claims implicit in option prices", *Journal of Business*, 51, 621-652
- Bulow, J. and P. Klemper (1994), "Rational frenzies and crashes," *Journal of Political Economy* 102, 1-23
- Campa, J.M. and P.H. Chang (1995), "Testing the expectations hypothesis of the term structure of volatilities in foreign exchange options", *Journal of Finance*, 50, 529-547
- Campa, J.M., P.H. Chang and R.L. Reider (1997), "ERM bandwidths for EMU and after: evidence from foreign exchange options", *Economic Policy*, 24, 55-89
- Canina, L. and S. Figlewski (1993), "The informational content of implied volatility", *Review of Financial Studies*, 6, 3, 659-681

- Cason, T.N. (2000), "The opportunity for conspiracy in asset markets organized with dealer intermediaries", *Review of Financial Studies* 13, 385-416
- CBOE (1975), "Analysis of volume and price patterns in stocks underlying CBOE options from December 30, 1974 to April 30, 1975", *Mimeo.*, Chicago.
- CBOE (1976), "Analysis of volume and price patterns in stocks underlying CBOE options from December 31, 1975 to January 16, 1976", *Mimeo.*, Chicago
- Christensen, B.J. and N.R. Prabhala (1998), "The relation between implied and realized volatility", *Journal of Financial Economics*, 50, 125-150
- Clewlow, L. and C. Strickland (2001), "Valuing energy options in a one factor model fitted to forward prices", *working paper*, University of Sydney
- Conrad, J. (1989), "The price effects of option introduction", *Journal of Finance* 44, 487-498
- Corrado, C. and T. Su (1996), "Skewness and kurtosis in S&P index returns implied by option prices", *Journal of Financial Research*, 19, 175-192
- Cox, J., S. Ross and M. Rubinstein (1979), "Option pricing: a simplified approach", *Journal of Financial Economics*, 7, 229-264
- Damodaran, A., and C. Lim (1991), "The effects of option listing on the underlying stocks return processes", *Journal of Banking and Finance* 15, 647-664
- Day, T.E. and C.M. Lewis (1992), "Stock market volatility and informational content of stock index options", *Journal of Econometrics*, 52, 267-287
- Day, T.M. and C.M. Lewis (1993), "Forecasting futures market volatility", *Journal of Derivatives*, 1, 33-50
- Deng, S. (2000), "Stochastic models of energy commodity prices and their applications: mean reversion with jumps and spikes", *working paper*, University of California Energy Institute
- Derman, E., I. Kani and N. Chriss (1996), "Implied trinomial trees of the volatility smile", *Journal of Derivatives*, Summer, 7-22

- Derman, E. (1999), "Regimes of volatility", *Risk*, April, 55-59
- Deutsche Bunderbank (1995), "The information content of derivatives for monetary policy", *Monthly report*, 47, 17-32
- Detemple, J., and P. Jorion (1990), "Option listing and stock returns", *Journal of Banking and Finance* 14, 781-801
- Dimson, E. and M. Marsh (1990), "Volatility forecasting without data-snooping", *Journal of Banking and Finance* 14, 399-421
- Ding, Z., C.W.J. Granger and R. Engle (1993), "A long memory property of stock market returns and a new model", *Journal of Empirical Finance*, 1, 83-106
- Duan, J.-C. (1995), "The GARCH option pricing model", *Mathematical Finance*, 5, 13-32
- Duan, J.-C. (1996), "Cracking the smile", *Riak*, 9, 55-59
- Duan, J.-C. and J.-G. Simonato (1998a), "Empirical martingale simulation for asset prices", *Management Science*, 44, 1218-1233
- Duan, J.-C. and J.-G. Simonato (1998b), "American option pricing under GARCH by a Markov chain approximation", Department of Finance, Hong Kong University of Science and Technology, *Working Paper*
- Duan, J.-C., G. Gauthier and J.-G. Simonato (1998), "An analytical approximation for the GARCH option pricing model", Department of Finance, Hong Kong University of Science and Technology, *Working Paper*
- Duan, J.-C. (1999), "Conditional Fat-Tailed Distributions and the Volatility Smile in Options", Department of Finance, Hong Kong University of Science and Technology, *Working Paper*
- Dumas, B., J. Fleming and R.E. Whaley (1998), "Implied volatility functions: empirical tests", *The Journal of Finance*, 53, 2059-2106

Easley, D. and M. O'Hara (1987), "Prices, trade size and information in security markets", *Journal of Financial Economics*, 19, 69-90

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

Engle, R.F. and V. Ng (1993), "Measuring and testing the impact of news on volatility", *Journal of Finance*, 48, 1749-78

Fedenia, M., and T. Grammatikos (1992), "Option trading and the bid-ask spreads of the underlying stocks", *Journal of Business* 65, 335-351

Federal Reserve Bank of Atlanta (Abken, P.A., 1995), "Using eurodollar futures options: gauging the market's view of interest rate movements", *Economic Review*, 16, 10-30

Fernandez, C. and M. Steel (1998), "On Bayesian modeling of fat tails and skewness", *Journal of the American Statistical Association*, 93, 359-371

Figlewski, S. (1997), "Forecasting volatility", *Financial Markets, Institutions and Instruments*, New York University Salomon Center, 6, 1, 1-88

Fleming, J., B. Ostdiek and R.E. Whaley (1995), "Predicting stock market volatility: a new measure", *Journal of Futures Markets*, 15, 229-235

Fleming, J., B. Ostdiek and R. Whaley (1996), "Trading costs and the relative rates of price discovery in stock, futures, and option markets", *Journal of Futures Markets*, 353-387

Fleming, J. (1998), "The quality of market volatility forecasts implied by S&P100 index option prices", *Journal of Empirical Finance*, 5, 317-345

Fleming, J. and B. Ostdiek (1999), "The impact of energy derivatives on the crude oil market", *Energy Economics* 21, 135-167

Fleten, S.E. and J. Lemming (2001), "Constructing forward prices in electricity markets", *working paper*, Riso National Laboratory Denmark

- Forsythe, R., Thomas R. Palfrey and C.R. Plott (1982), "Asset valuation in an experimental market", *Econometrica* 50, 537-567
- Forsythe, R., T.R. Palfrey and C.R. Plott (1984), "Futures markets and informational efficiency: A laboratory examination", *Journal of Finance* 39, 955-981
- Franses, P.H. and D. van Dijk (1995), "Forecasting stock market volatility using (non-linear) GARCH models", *Journal of Forecasting*, 15, 229-235
- Friedman, Daniel, Glenn Harrison, and J. Salmon (1984), "The informational efficiency of experimental asset markets", *Journal of Political Economy* 92, 349-408
- Friedman, D. and S. Sunder (1994), "Experimental methods: a primer for economists", Cambridge University Press
- Glosten, L.R., R. Jagannathan and D.E. Runkle (1993), "On the relation between the expected value and the volatility of the nominal excess return on stocks", *Journal of Finance*, 48, 1779-1801
- Granger, C.W.J. (1969), "Investigating Causal Relations by Econometric Models and Cross-spectral Models", *Econometrica*, 37, 424-438
- Hafner, C.M. and H. Herwartz (2001), "Option pricing under linear autoregressive dynamics, heteroskedasticity, and conditional leptokurtosis", *Journal of Empirical Finance*, 8, 1-34
- Hamilton, J.D. (1989), "A new approach to the economic analysis of non-stationary time series and the business cycle", *Econometrica*, 57, 357-384
- Hamilton, J.D. (1994), "Time Series Analysis", Princeton University Press
- Hanke, M. (1997), "Neural network approximation of option pricing formulas for analytically intractable option pricing models", *Journal of Computational Intelligence in Finance*, 5, 20-27
- Hansen, B. (1994), "Autoregressive conditional density estimation", *International Economic Review*, 35, 705-730

Harvey, A.C. (1989), "Forecasting, structural time series and the Kalman filter", Cambridge University Press, Cambridge

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

Heston, S.L. and S. Nandi (2000), "A closed-form GARCH option valuation model", *Review of Financial Studies*, 3, 585-625

Heynen, R., A. Kemma and T. Vorst (1994), "Analysis of the term structure of implied volatilities", *Journal of Financial and Quantitative Analysis*, 29, 31-56

Hilliard, J.E. and J. Reis (1998), "Valuation of commodity futures and options under stochastic convenience yields, interest rates, and jump diffusions in the spot", *Journal of Financial and Quantitative Analysis*, 33:1, 61-86

Ho, T.S.Y. and S. Lee (1986), "Term structure movements and pricing interest rate contingent claims", *Journal of Finance*, 41, 1011-1029

Huisman, R. and R.J. Mahieu (2001), "Regime jumps in power prices", *Energy & Power Risk Management*, September

Hull, J. and A. White (1987), "The pricing of options on assets with stochastic volatilities", *Journal of Finance* 42, 281-300

Jackwerth, J. and M. Rubinstein (1996), "Recovering probability distributions from option prices", *Journal of Finance*, 5, 1611-1631

Jackwerth, J. (1997), "Generalized binomial trees", *Journal of Derivatives*, Winter, 7-17

Jackwerth, J. (2000), "Recovering risk aversion from option prices and realized returns", *Review of Financial Studies*, 13, 433-467

Jarrow, R. and A. Rudd (1982), "Approximate option valuation for arbitrary stochastic processes", *Journal of Financial Economics*, 10, 347-369

- Jennings, R.H., and L.T. Starks (1986), "Earnings announcements, stock price adjustment, and the existence of option Markets", *Journal of Finance* 41, 107-125
- John, K., A. Koticha, R. Narayan and M. Subrahmanyam (2000), "Margin Rules, Informed Trading in Derivatives, and Price Dynamics", *working paper*
- Johnson and Barz (1999), "Selecting stochastic processes for modeling electricity prices", *Energy Modeling and the Management of Uncertainty*, Risk Publications
- Jorion, P. (1995), "Predicting volatility in the foreign exchange market", *Journal of Finance*, 50, 507-528
- Kagel, J.H. and A.E. Roth (1995), "The handbook of experimental economics", Princeton University Press
- Kallsen, J. and M. Taqqu (1998), "Option pricing in ARCH-type models", *Mathematical Finance*, 8, 13-26
- Kaminski, V. (1997), "The challenge of pricing and risk managing electricity derivatives", *The US Power Market*, Risk Publications, 149-171
- Klemkosky, R. and T. Maness (1980), "The impact of options on the underlying securities", *Journal of Portfolio Management* 6, 12-18
- Kluger, B.D., and S.B. Wyatt (1995), "Options and efficiency: some experimental evidence", *Review of Quantitative Finance and Accounting* 5, 179-201
- Knittel, C.R. and M. Roberts (2001), "An empirical examination of deregulated electricity prices", *working paper*, University of California Energy Institute
- Koekebakker, S. and F. Ollmar (2001), "Forward curve dynamics in the Nordic electricity market", *working paper*, Agder University College, Denmark
- Kumar, R., A. Sarin, and K. Shastri (1998), "The impact of options trading on the market quality of the underlying security: An empirical analysis", *Journal of Finance* 53, 717-732
- Kyle, A. S. (1985), "Continuous auctions and insider trading", *Econometrica* 53, 1315-1335

- Lambert, P. and S. Laurent (2001), "Modeling financial time-series using GARCH-type models and a skewed Student density", *mimeo*, Université de Liège
- Lamoureux, C.G. and C.R. Schnitzlein (1997), "When it's not the only game in town: The effect of bilateral search on the quality of a dealer market", *Journal of Finance*, 52, 683-712
- Leahy, M.P. and C.P. Thomas (1996), "The sovereignty option: The Quebec referendum and market views on the Canadian Dollar", *Board of Governors of the Federal Reserve System International Finance Discussion Papers*, 555
- Li, K. (2002), "Long memory versus option-implied volatility predictions", *Journal of Derivatives*, 9, 9-25
- Longstaff, F.A. and E.S. Schwartz (2001), "Valuing American options by simulation: a simple least-square approach", *Review of Financial Studies*, 14:1, 113-147
- Longstaff, F.A., P. Santa-Clara and E.S. Schwartz (2001), "Throwing away a billion dollars: the cost of sub-optimal exercise strategies in the swaption market", *Journal of Financial Economics*, 62:1, 39-66
- Lucia, J. and E.S. Schwartz (2002), "Electricity prices and power derivatives: evidence from the Nordic Power Exchange", *Review of Derivatives Research* 5, 5-50
- Madan, D. and F. Milne (1991), "Option pricing with VG martingale components", *Mathematical Finance*, 1, 39-56
- Malz, A. (1996), "Using option prices to estimate realignment probabilities in the European Monetary System: the case of Sterling-Mark", *Journal of International Money and Finance*, 15, 717-748
- Malz, A. (1997), "Estimating the probability distribution of the future exchange rate from option prices", *Journal of Derivatives*, 5, 18-36
- Martens, M. and J. Zein (2002), "Predicting financial market volatility: High-frequency time-series forecasts vis-à-vis implied volatility", *working paper*

Mayhew, S., and V. Mihov (2001), "Another look at option listing effects", *working paper*

Melick, W. and C. Thomas (1997), "Recovering an asset's implied PDF from option prices: an application to crude oil during the Gulf Crisis", *Journal of Financial and Quantitative Analysis*, March, 91-115

Merton, R.C. (1976), "Option-pricing when underlying stock returns are discontinuous", *Journal of Financial Economics*, 3, 125-144

Milgrom, P. and N. Stokey (1982), "Information, trade, and common knowledge", *Journal of Economic Theory*, 26, 17-27

Miltersen, K. and E.S. Schwartz (1998), "Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates", *Journal of Financial and Quantitative Analysis*, 33:1, 33-59

Mitnik, S. and M.S. Paoletta (2000), "Conditional density and value-at-risk prediction, of Asian currency exchange rates", *Journal of Forecasting*, 19, 313-333

Nagot, I. and R. Trommsdorff (1999), "The tree of knowledge", *Risk*, August, 99-102

Naik, V. and M.H. Lee (1990), "General equilibrium pricing of options on the market portfolio with discontinuous returns", *Review of Financial Studies*, 3, 493-522

Nelson, D. (1991), "Conditional heteroskedasticity in asset returns: a new approach", *Econometrica*, 59, 347-370

Nelson, D. (1992), "Filtering and forecasting with misspecified ARCH models I: getting the right variance with the wrong model", *Journal of Econometrics*, 52, 61-90

Nelson, D. and D.P. Foster (1995), "Filtering and forecasting with misspecified ARCH models II: making the right forecast with the wrong model", *Journal of Econometrics*, 67, 303-335

Pilipovic, D. (1998), "Energy risk: valuing and managing energy derivatives", McGraw Hill, New York

- Pindyck, R.S. (1999), "The long-run evolution of energy prices", *The Energy Journal*, 20:2, 1-27
- Pirrong, C. and M. Jermakyan (1999), "Valuing power and weather derivatives on a mesh using finite difference methods", *Energy Modeling and the Management of Uncertainty*, Risk Books
- Pirrong, C. and M. Jermakyan (2000), "The price of power: The valuation of power and weather derivatives", *working paper*, Washinton University
- Plott, C.R., and S. Sunder (1988), "Rational expectations and the aggregation of diverse information in laboratory security markets," *Econometrica* 56, 1085-1118
- Pong, S., M.B. Stackleton, S.J. Taylor and X. Xu (2002), "Forecasting Sterling/Dollar volatility: implied volatility versus long memory intraday models", *Working paper* Lancaster University
- Reiss, R.D. and M. Thomas (1997), "Statistical analysis of extreme values: with applications to insurance, finance, hydrology, and other fields", 2nd ed., Birkhauser, Basel
- Ritchken, P. and R. Trevor (1999), "Pricing options under generalized GARCH and stochastic volatility processes", *Journal of Finance*, 54, 377-402
- Rubinstein, M. (1994), "Implied binomial trees", *Journal of Finance*, 49, 771-818
- Schwartz, E.S. (1997), "The stochastic behavior of commodity prices: implications for valuation and hedging", *Journal of Finance* 52:3, 923-973
- Sheikh, A.M., and E.I. Ronn (1994), "A characterization of the daily and intraday behavior of returns on options", *Journal of Finance* 49, 557-580
- Shimko, D. (1993), "Bounds of probability", *Risk*, April, 33-37
- Skinner, D.J. (1989), "Options markets and stock return volatility", *Journal of Financial Economics* 23, 61-78
- Stein, E. and J. Stein (1991), "Stock price distributions with stochastic volatility: an analytic approach", *Review of Financial Studies*, 4, 727-752

Trennepohl, G.L., and W. Dukes (1979), "CBOE options and stock volatility", *Review of Business and Economic Research* 14, 49-60

Vijh, A.M. (1990), "Liquidity of the CBOE equity options", *Journal of Finance*, 45, 1157-1180

Whiteside, M.M., W. P. Dukes, and Patrick M. Dunne (1983), "Short-term impact of option trading on underlying securities", *Journal of Financial Research* 6, 313-321

Wiggins, J. (1987), "Option values under stochastic volatility: theory and empirical estimates", *Journal of Financial Economics*, 19, 351-372

Samenvatting (Summary in Dutch)

Dit proefschrift beschrijft vijf onderzoeken op het gebied van financiële derivaten. Hier volgt een samenvatting.

deel I: de microstructuur van derivatenmarkten

In het eerste deel van dit proefschrift onderzoeken we de invloed van de handel in derivaten (zoals opties) op de handel in het onderliggende waardepapier (zoals aandelen). We toetsen in het bijzonder de hypothese dat een derivaat de efficiëntie in het onderliggende waardepapier verbetert. De twee gepresenteerde studies ondersteunen grotendeels deze hypothese en verduidelijken de mechanismen die hiertoe leiden. Een plausibele verklaring is dat een gecorreleerd waardepapier ervoor zorgt dat prijsontwikkeling op twee markten tegelijkertijd plaatsvindt. Market makers in het onderliggende waardepapier kunnen betere prijzen afgeven als ze leren van de ontwikkelingen in de optiemarkt.

In hoofdstuk 3 maken we gebruik van experimenten, een gecontroleerde handelsomgeving waarin studenten handelen in een aandeel en een call optie op dat aandeel. De handel vindt plaats onder asymmetrische informatie, wat concreet betekent dat maar een handelaar (de insider) de werkelijke waarde van het aandeel weet. De experimenten maken het mogelijk om alle informatie en alle handelingen nauwkeurig te observeren in een opzet die lijkt op het theoretische model van Kyle (1985), maar veel realistischer is dan enig theoretisch model. Herhaalde handelsrondes met verschillende groepen studenten en verschillende aandeelwaarden maken duidelijk dat een insider actief in beide markten handelt, met een voorkeur voor de markt die op dat moment de meeste winstpotentie biedt. Beide markten hebben zo hun aandeel in het naar de oppervlakte komen van de werkelijke waarde van aandeel en optie. Tussen beide markten vindt namelijk een continue interactie plaats: transacties in de ene markt leiden tot aanpassingen in de geboden en gevraagde prijzen in de andere markt.

Doordat een vergelijking tussen een markt met en een zonder opties leidt tot controleproblemen (vanwege een verschillend aantal marktdeelnemers), hebben we

ervoor gekozen om de invloed van opties indirect te bestuderen. Dit bereiken we door onderscheid te maken tussen handelsrondes waarin de optie geen waarde heeft en handelsrondes waarin de optie wel een waarde heeft. De resultaten laten zien dat dit onderscheid een grote invloed heeft op de prijzen in de aandelenmarkt: als de optie in-the-money is (waarde heeft), convergeren de prijzen veel sneller en met minder volatiliteit naar hun werkelijke waarde. De neiging van insiders om daar te handelen waar de grootste winst te behalen is, zorgt voor een rijker scala aan signalen waarop de market makers hun prijzen kunnen baseren, dan wanneer de insider in maar een markt winstgevend kan handelen. Daarmee laten we zien dat een optie niet alleen het prijsontwikkelingsproces in tweeën splitst (aandeel is meer of minder waard dan de uitoefenprijs van de optie), maar ook het proces verandert waarmee de marktdeelnemers hun verwachtingen aanpassen. We laten bovendien zien dat dit effect sterker is bij een minder strategisch handelende insider (vanwege risico-aversie, ongeduld of slechte informatie bijvoorbeeld).

Het model in hoofdstuk 4 biedt twee belangrijke toevoegingen aan de bestaande theoretische modellen over de interactie tussen aandelen en opties. Ten eerste breidt het de bestaande enkelvoudige handelsmodellen uit tot een meervoudig handelsmodel (waarbij meerdere transacties kunnen plaatsvinden). Ten tweede analyseren we met het model verschillende hefboorniveaus van opties; het hefboomeffect is immers een belangrijke onderscheidende eigenschap van opties ten opzichte van aandelen. We beginnen met een standaard sequentieel handelsmodel en laten zien dat het dynamisch is. Doordat handelaren hun verwachtingen na iedere transactie aanpassen, kunnen we een reeks transacties simuleren en preciezere maatstaven voor marktefficiëntie bestuderen. Het model verduidelijkt en scheidt twee mechanismen die volgen op het introduceren van een optie. Enerzijds leren handelaren en market makers van de transacties in de optiemarkt en kunnen nauwkeurigere prijzen afgeven. Anderzijds beïnvloedt de introductie van een optie de samenstelling van de handelaren: het aantal insiders verandert het aantal insiders, wat mogelijk leidt tot een slechter functionerende markt, afhankelijk van de gebruikte maatstaf. De resultaten van ons dynamische model duiden op een afname in prijsfouten (verschil tussen werkelijke waarde en handelsprijs) onder alle onderzochte parameterinstellingen. Ongeïnformeerde handelaren profiteren hiervan het meest in een markt waarin grote transacties kunnen worden uitgevoerd en in een markt met verhoudingsgewijs weinig insiders. Dit komt overeen met een goedontwikkelde derivatenmarkt waarin opties voor een belangrijk hefboomeffect zorgen. Dezelfde

kenmerken van een goed ontwikkelde derivatenmarkt hebben op de volatiliteit van het aandeel juist een tegengesteld (negatief) effect: de volatiliteit neemt toe.

Deel II: Empirische studies in derivatenmarkten

In het tweede deel van dit proefschrift komen drie empirische derivatenstudies aan bod die alle uitgaan van niet-normaal verdeelde rendementen. In de eerste twee studies onderzoeken we methoden om informatie te af te leiden uit marktprijzen van opties. Het laatste hoofdstuk behandelt een traditioneel optiewaarderingsonderzoek, maar toegepast op een ongewoon en grillig onderliggend goed, namelijk elektriciteit.

Hoofdstuk 7 behandelt een flexibele methode om de *risico-neutrale* verdeling van het onderliggende waardepapier uit optieprijsen te herleiden. De exacte vorm van deze geïmpliceerde verdeling geeft belangrijke informatie die ingezet kan worden om andere opties op het onderliggende waardepapier te waarderen, om vergelijkingen te maken tussen opties op verschillende onderliggende waardepapieren, en om nauwgezet het sentiment in de markt te volgen. We passen een scheve variant van de welbekende Student-t verdeling toe om de veelvoorkomende vormen ('smiles' en 'skews') in de de geïmpliceerde volatiliteitcurve te modelleren. De belangrijkste kracht van deze methode is de directe parametrisering van scheefheid en dikstaartigheid in de geïmpliceerde verdeling. In een toepassing op verschillende jaren van FTSE-100 index opties vergelijken we de geschiktheid van deze methode met de normaal-verdelingsmethode (constante volatiliteit), twee curve-methoden (die de geïmpliceerde volatiliteitcurve modelleren en daaruit risico-neutrale verdelingen afleiden) en een trinomiaale boomstructuur.

Van de onderzochte methoden levert de (standaard) normaal-verdelingsmethode veruit het slechtste resultaat. Van de overige methoden levert de curve-methode (met optiedelta als verklarende variabele voor geïmpliceerde volatiliteit) het beste resultaat: optieprijsen uit dit model komen het sterkst overeen met werkelijke prijzen en dit effect bestendigt in latere jaren. Hoewel de scheve Student-t methode en de trinomiaale boom een minder goede benadering geven van werkelijke prijzen, hebben ze beide belangrijke toepassingsmogelijkheden. De trinomiaale boomstructuur is geschikt om Amerikaanse-type opties te waarderen, terwijl parameters van de scheve Student-t methode direct gerelateerd zijn aan de momenten van de onderliggende verdeling, waardoor veranderingen in marktsentiment goed gevolgd kunnen worden..

Hoofdstuk 8 beschrijft een methode om uit optiepreizen een voorspelling af te leiden van het werkelijke prijsproces van het onderliggende waardepapier, in plaats van slechts een verdeling op een enkel moment in de toekomst (zoals de methode in hoofdstuk 7). Aan de hand van het optiewaarderingsmodel van Duan (1995) bepalen we een geïmpliceerd EGARCH-proces en op basis daarvan een eendaagse voorspelling van de volatiliteit van de onderliggende index. We vergelijken deze voorspelling vervolgens met een voorspelling die geheel gebaseerd is op historische rendementen. Dit doen we voor de Engelse FTSE-100 en de Duitse DAX-30 index, waar ze tot verschillende resultaten leiden. Op de Duitse markt lopen de voorspellingen van het geïmpliceerde proces vooruit op de voorspellingen uit historische rendementen, maar dit is niet het geval op de Engelse markt. In de Duitse markt verklaren de historische rendementen ongeveer 23% van de werkelijke gerealiseerde volatiliteit (die we construeren met intra-dag rendementen), terwijl de geïmpliceerde voorspelling anderhalf keer meer verklaart. In de Engelse markt verklaren de historische en geïmpliceerde voorspelling beide niet minder dan 44% van de gerealiseerde volatiliteit. Hier bovenop verklaart een gewogen gemiddelde voorspelling in de Engelse markt zelfs nog 8% extra. In de Duitse markt daarentegen bevat de geïmpliceerde voorspelling nagenoeg alle informatie over de toekomst. Tenslotte is het vermeldenswaard dat de geïmpliceerde voorspelling in de Duitse markt nauwelijks vertekening vertoont in het voorspellen van gerealiseerde volatiliteit (inclusief nachtrendementen), maar wel vertekent in de Engelse markt. Samenvattend kunnen we concluderen dat de geïmpliceerde voorspelling in Engeland een goede voorspeller is van gerealiseerde volatiliteit, maar niet beter is dan een historische voorspelling en bovendien vertekend is. In de Duitse markt verdient de geïmpliceerde voorspelling duidelijk de voorkeur.

In hoofdstuk 9 beschrijven we een zogenaamd regime-switch model om opties te waarderen op spot elektriciteitsprijzen (prijzen voor levering een dag nadien). Het model bevat de belangrijkste eigenschappen van spot elektriciteitsprijzen: seizoensafhankelijkheid, middelpunt-tenderend (mean-reverting), hoge volatiliteit en plotselinge pieken. Een belangrijke bijdrage van het model is dat het meerdere pieken na elkaar toelaat, wat belangrijk is voor risico-management toepassingen en waardering van derivaten. Uit dit spot prijsmodel leiden we formules af voor de waardering van Europese-stijl call en put opties. Doordat het regime-switch model bestaat uit twee onafhankelijke regimes kunnen we optiepreizen in twee componenten opsplitsen: een voor het stabiele middelpunt-tenderende proces, en een de pieken.

Beide componenten waarderen we met het model van Black (1976). In een toepassing op Nederlandse APX spot prijzen verduidelijken we het belang van een juiste modellering van pieken voor de waardering van opties: optiepreizen van het regime-switch model zijn beduidend hoger en komen beter overeen met historische uitbetalingen dan prijzen van een model zonder pieken. Dit effect is het duidelijkst voor de veelverhandelde out-of-the-money call opties.

Conclusie

Beide studies in het eerste deel van dit proefschrift verduidelijken de complexe relatie en afhankelijkheid tussen twee gecorreleerde waardepapieren, een aandeel en een optie. De experimenten tonen aan dat een optie ervoor zorgt dat de efficiëntie in de markt verbetert, maar het theoretische model plaatst hier de kanttekening bij dat dit resultaat gedeeltelijk af kan hangen van de karakteristieken van de markt. Zonder rekening te houden met natuurlijk menselijk gedrag onderschat het theoretisch model echter enigszins de zegeningen van een optiemarkt. In het model trekken bijvoorbeeld alle deelnemers aan een markt dezelfde conclusies uit iedere transactie. Daaruit volgt dat er geen verschillen van inzicht bestaan tussen market makers in de aandelenmarkt en die in de optiemarkt. In de experimenten, net als in de praktijk, zijn die verschillen (en de reactie hierop van insiders) juist erg informatief en bespoedigen een juiste prijsontwikkeling.

In het tweede deel van dit proefschrift bouwen we voort op de positieve conclusie van het eerste gedeelte en presenteren drie optiewaarderingsmethoden. Hoewel het vroege werk van Black, Scholes en Merton de ontwikkeling van derivatenmarkten enorm gestimuleerd heeft, laten we zien dat de veronderstellingen in hun modellen aangepast dienen te worden voor realistischere optieprijswaarderingen. Op basis van een recente innovatie herleiden we bijvoorbeeld uit marktprijzen van opties een prijsproces dat belangrijke informatie kan bevatten over de toekomstige ontwikkeling van een marktindex. Voor de zeer specifieke elektriciteitsmarkt ontwikkelen we bovendien een model om opties op spotprijzen adequaat te waarderen. Met de ontwikkeling van dergelijke modellen beogen we derivatenmarkten beter te laten functioneren, wat leidt tot een groter vertrouwen in deze markten en een betere spreiding van risico's onder marktdeelnemers.

Curriculum vitae

Cyriel de Jong was born on 19 July 1976 in Valkenburg aan de Geul, the Netherlands. From 1994 till 1999 he studied econometrics at Maastricht University. During this time he spent half a year as an exchange student at the Universität Wien in Vienna. He furthermore completed an internship on mortgage prepayment modelling at De Nationale Investeringsbank in The Hague and a research project on European Corporate Bonds at Maastricht University. He obtained his MSc in econometrics in 1999 with honour.

From February 1999 onwards Cyriel de Jong has been a PhD student at the Financial Management department of the Rotterdam School of Management at Erasmus University. During this period he has taught several executive and non-executive courses. From 2001 till 2003 he furthermore worked for the consultancy firm FinEdge, where he was responsible for the subsidiary Energy Global. His research was published in Energy Power Risk Management, Bedrijfskunde and various Dutch journals. Since 1 May 2003 he has been assistant professor at the Financial Management department at Erasmus University, where his research interests focus on commodity markets in general and derivative valuation in particular. Besides his academic career he continues working part-time as a consultant in financial and commodity markets.

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT (ERIM)

ERIM PH.D. SERIES
RESEARCH IN MANAGEMENT

ERIM Electronic Series Portal: <http://hdl.handle.net/1765/1>

Title: **Operational Control of Internal Transport**
Author: J. Robert van der Meer
Promotor(es): Prof.dr. M.B.M. de Koster, Prof.dr.ir. R. Dekker
Defended: September 28, 2000
Series number: 1
ISBN: 90-5892-004-6

Title: **Quantitative Models for Reverse Logistics**
Author: Moritz Fleischmann
Promotor(es): Prof.dr.ir. J.A.E.E. van Nunen, Prof.dr.ir. R. Dekker, dr. R. Kuik
Defended: October 5, 2000
Series number: 2
Published: Lecture Notes in Economics and Mathematical Systems,
Volume 501, 2001, Springer Verlag, Berlin,
3540 417 117
ISBN: 3540 417 117

Title: **Optimization Problems in Supply Chain Management**
Author: Dolores Romero Morales
Promotor(es): Prof.dr.ir. J.A.E.E. van Nunen, dr. H.E. Romeijn
Defended: October 12, 2000
Series number: 3
ISBN: 90-9014078-6

Title: **Layout and Routing Methods for Warehouses**
Author: Kees Jan Roodbergen
Promotor(es): Prof.dr. M.B.M. de Koster, Prof.dr.ir. J.A.E.E. van Nunen
Defended: May 10, 2001
Series number: 4
ISBN: 90-5892-005-4

Title: **Rethinking Risk in International Financial Markets**
Author: Rachel Campbell
Promotor(es): Prof.dr. C.G. Koedijk
Defended: September 7, 2001
Series number: 5
ISBN: 90-5892-008-9

Title: **Labour flexibility in China's companies: an empirical study**
Author: Yongping Chen
Promotor(es): Prof.dr. A. Buitendam, Prof.dr. B. Krug
Defended: October 4, 2001
Series number: 6
ISBN: 90-5892-012-7

Title: **Strategic Issues Management: Implications for Corporate Performance**
Author: Pursey P.M.A.R. Heugens
Promotor(es): Prof.dr.ing. F.A.J. van den Bosch, Prof.dr. C.B.M. van Riel
Defended: October 19, 2001
Series number: 7
ISBN: 90-5892-009-7

Title: **Beyond Generics; A closer look at Hybrid and Hierarchical Governance**
Author: Roland F. Speklé
Promotor(es): Prof.dr. M.A. van Hoepen RA
Defended: October 25, 2001
Series number: 8
ISBN: 90-5892-011-9

Title: **Interorganizational Trust in Business to Business E-Commerce**
Author: Pauline Puvanasvari Ratnasingam
Promotor(es): Prof.dr. K. Kumar, Prof.dr. H.G. van Dissel
Defended: November 22, 2001
Series number: 9
ISBN: 90-5892-017-8

Title: **Outsourcing, Supplier-relations and Internationalisation:
Global Source Strategy as a Chinese puzzle**

Author: Michael M. Mol
Promotor(es): Prof.dr. R.J.M. van Tulder
Defended: December 13, 2001
Series number: 10
ISBN: 90-5892-014-3

Title: **The Business of Modularity and the Modularity of Business**

Author: Matthijs J.J. Wolters
Promotor(es): Prof.dr. P.H.M. Vervest, Prof.dr.ir. H.W.G.M. van Heck
Defended: February 8, 2002
Series number: 11
ISBN: 90-5892-020-8

Title: **The Quest for Legitimacy; On Authority and Responsibility in
Governance**

Author: J. van Oosterhout
Promotor(es): Prof.dr. T. van Willigenburg, Prof.dr. H.R. van Gunsteren
Defended: May 2, 2002
Series number: 12
ISBN: 90-5892-022-4

Title: **Information Architecture and Electronic Market Performance**

Author: Otto R. Koppius
Promotor(es): Prof.dr. P.H.M. Vervest, Prof.dr.ir. H.W.G.M. van Heck
Defended: May 16, 2002
Series number: 13
ISBN: 90-5892-023 - 2

Title: **Planning and Control Concepts for Material Handling Systems**

Author: Iris F.A. Vis
Promotor(es): Prof.dr. M.B.M. de Koster, Prof.dr.ir. R. Dekker
Defended: May 17, 2002
Series number: 14
ISBN: 90-5892-021-6

Title: **Essays on Agricultural Co-operatives;
Governance Structure in Fruit and Vegetable Chains**
Author: Jos Bijman
Promotor(es): Prof.dr. G.W.J. Hendrikse
Defended: June 13, 2002
Series number: 15
ISBN: 90-5892-024-0

Title: **Analysis of Sales Promotion Effects on Household Purchase
Behavior**
Author: Linda H. Teunter
Promotor(es): Prof.dr.ir. B. Wierenga, Prof.dr. T. Kloek
Defended: September 19, 2002
Series number: 16
ISBN: 90-5892-029-1

Title: **Incongruity between Ads and Consumer Expectations of
Advertising**
Author: Joost Loef
Promotor(es): Prof.dr. W.F. van Raaij, Prof.dr. G. Antonides
Defended: September 26, 2002
Series number: 17
ISBN: 90-5892-028-3

Title: **Creating Trust between Local and Global Systems**
Author: Andrea Ganzaroli
Promotor(es): Prof.dr. K. Kumar, Prof.dr. R.M. Lee
Defended: October 10, 2002
Series number: 18
ISBN: 90-5892-031-3

Title: **Coordination and Control of Globally Distributed Software
Projects**
Author: Paul C. van Fenema
Promotor(es): Prof.dr. K. Kumar
Defended: October 10, 2002
Series number: 19
ISBN: 90-5892-030-5

Title: **Improving the flexibility and profitability of ICT-enabled business networks: an assessment method and tool.**
Author: Dominique J.E. Delporte- Vermeiren
Promotor(es): Prof.mr.dr. P.H.M. Vervest, Prof.dr.ir. H.W.G.M. van Heck
Defended: May 9, 2003
Series number: 20
ISBN: 90-5892-040-2

Title: **Organizing Knowledge in Internal Networks. A Multilevel Study**
Author: Raymond van Wijk
Promotor(es): Prof.dr.ing. F.A.J. van den Bosch
Defended: May 22, 2003
Series number: 21
ISBN: 90-5892-039-9

Title: **Cyclic Railway Timetable Optimization**
Author: Leon W.P. Peeters
Promotor(es): Prof.dr. L.G. Kroon, Prof.dr.ir. J.A.E.E. van Nunen
Defended: June 6, 2003
Series number: 22
ISBN: 90-5892-039-9

Dealing with Derivatives:

Studies on the role, informational content and pricing of financial derivatives

The aim of this thesis is to improve the understanding of derivatives markets, which should ultimately lead to a better diversification of risks among market participants. The author first analyzes the impact of derivatives on the market quality of the underlying asset. With experiments and a theoretical model it is shown that derivatives generally make markets more efficient, although volatility may increase, depending on the exact market structure. Next, the author presents two methods that derive information about the underlying price process from traded options. The models approximate the option prices well and the extracted information explains future volatility better than historical data. Finally, a model for the valuation of options in electricity markets is presented that deals with the special characteristics of electricity spot prices and may serve to value electricity generation plants.

ERIM

The Erasmus Research Institute of Management (ERIM) is the Research School (Onderzoekschool) in the field of management of the Erasmus University Rotterdam. The founding participants of ERIM are the Rotterdam School of Management and the Rotterdam School of Economics. ERIM was founded in 1999 and is officially accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW). The research undertaken by ERIM is focussed on the management of the firm in its environment, its intra- and inter-firm relations, and its business processes in their interdependent connections. The objective of ERIM is to carry out first rate research in management, and to offer an advanced graduate program in Research in Management. Within ERIM, over two hundred senior researchers and Ph.D. candidates are active in the different research programs. From a variety of academic backgrounds and expertises, the ERIM community is united in striving for excellence and working at the forefront of creating new business knowledge.

The ERIM Ph.D Series contains Dissertations in the field of Research in Management defended at Erasmus University Rotterdam. The Dissertations in the Series are available in two ways, printed and electronic. ERIM Electronic Series Portal: www.research-in-management.nl/