PhD Thesis

Empirical Essays on Historical Volatility Models, Option-Implied Volatility and the Efficiency of Options Markets

Author: Nikolaos Voukelatos

Supervisor: Professor Mark Shackleton

This thesis is submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

University of Lancaster

Department of Accounting & Finance

July 2009

Abstract

This thesis consists of four empirical essays on historical volatility models, implied volatility and the efficiency of options markets. The first essay examines the 'asymmetric volatility' phenomenon in index returns from the perspective of the 'diversification effect'. Changes in the average realized correlation among the Dow Jones's constituents are found to drive changes in the index's realized volatility and to be negatively and asymmetrically correlated with index returns. In line with the 'diversification hypothesis', it is shown that accounting for correlation dynamics in a GARCH specification reduces the level of asymmetry in index returns, with conditional correlation changes absorbing part of the past returns' explanatory power over changes in the conditional variance, while the returns' sign remains highly significant. The second essay examines volatility asymmetry from the perspective of the 'down-market effect', and reports that individual stocks' volatilities respond asymmetrically to lagged market returns and that the respective degree of asymmetry is comparable to the one exhibited by the index. The third essay explores the possibility of an option-implied measure of the exchange rate's future variance having an impact on the validity of the Uncovered Interest Parity condition. The results suggest that accounting for the above Jensen's Inequality Term significantly increases the proportion of forward unbiasedness regressions that are in line with theoretical parity predictions. Finally, the fourth essay focuses on the efficiency of the emerging Greek options market in terms of options' returns that are commensurate with the underlying risk. Based on a set of commonly used market efficiency measures, the hypothesis that the developing Greek options market exhibits a degree of efficiency that is comparable to those of the developed US and UK markets cannot be rejected.

Acknowledgements

I would like to express my gratitude to my supervisor Professor Mark Shackleton for giving me the opportunity to work with him in the PhD programme as well as for his guidance and encouragement throughout the whole process. This thesis has greatly benefited by his constructive comments and his constant support is gratefully acknowledged. I am also grateful to Dr Martin Widdicks for valuable comments on the fifth Chapter, to Efthymios Pavlidis for his assistance with the FM-LAD estimator in the fourth Chapter, and to Leonidas Tsiaras for our helpful discussion on the 'diversification effect' in the second Chapter.

I would like to thank my family for their unconditional support over the years. Their understanding and constant encouragement have had a significant impact on my studies and a large part of my motivation stems from them. Finally, I dedicate this thesis to Alexandra for her patience and support throughout the project. Her love, advice and enthusiasm were fundamental in all the stages of my PhD.

Declaration

I hereby declare that the contents of this thesis have not been previously submitted for the award of a higher degree in any university. To the best of my knowledge, this thesis contains no material previously published or written by any other person except where due references have been made in the thesis.

Nikolaos Voukelatos

March 2009

Table of Contents

Abstract			**i**
Acknowledgements			**ii**
Declaration			**iii**
Table of Contents			**iv**
Frequently Used Notation			**viii**
List of Acronyms and Abbreviations			**ix**
List of Figures			**x**
List of Tables			**xi**
Chapter 1	**Introduction**		**1**
Chapter 2	**The Role of Realized Correlation Dynamics in Explaining Volatility Asymmetry in Dow Jones Index Returns**		**10**
2.1	Introduction		10
	2.1.1	Literature Review	10
	2.1.2	Scope of Study	14
2.2	Data		16
2.3	Estimation of the Realized Volatility and the Average Realized Correlation of the DJIA		19
2.4	Volatility Asymmetry in DJIA Index Returns		22
	2.4.1	Daily Returns of the DJIA and of its Constituents	22
	2.4.2	Volatility Asymmetry in Index and in Individual Stocks' Returns	26
	2.4.3	Asymmetric GARCH Model	29
2.5	Analysis of the DJIA Average Realized Correlation		33
	2.5.1	Time-Series Properties	33
	2.5.2	Correlation Persistence	35
	2.5.3	Co-movement of Correlation and Volatility	38

	2.5.4	Correlation Asymmetry	41
2.6	The Impact of Correlation Dynamics on DJIA Volatility Asymmetry		43
	2.6.1	Extended Asymmetric Regressions	43
	2.6.2	Extended GJR-C Model	45
2.7	Conclusion		50
	References		53
Chapter 3	**The Asymmetric Impact of Firm-Specific and of Index Returns on the Volatility Processes of Individual Stocks**		**57**
3.1	Introduction		57
	3.1.1	Literature Review	57
	3.1.2	Scope of Study	59
3.2	Data		60
3.3	Volatility Asymmetry with respect to Firm-specific Returns		63
3.4	Volatility Asymmetry with respect to Market Returns		66
3.5	Volatility Asymmetry with respect to both Market and Stock Returns		69
3.6	Conclusion		70
	References		72
Chapter 4	**An Examination of the Forward Premium Puzzle Using Option Implied Information**		**74**
4.1	Introduction		74
	4.1.1	Literature Review	74
	4.1.2	Scope of Study	77
4.2	Economics		78
4.3	Data		81
	4.3.1	Sources	81
	4.3.2	Time-Series Properties	82
4.4	Estimating Option-Implied Variances		88
	4.4.1	Black and Scholes ATM Implied Variance	88
	4.4.2	Corrado and Su Implied Variance	91

4.4.3 Model-Free Implied Variance 92
4.5 OLS Estimation 97
4.5.1 Estimating the Standard Specification 97
4.5.2 Accounting for the Future Variance of the Spot Rate 99
4.6 FM-LAD Estimation 102
4.6.1 Complications Arising from Using OLS Regressions 102
4.6.2 The FM-LAD Estimator 103
4.6.3 FM-LAD Results 106
4.6.3.1 The Standard Specification 106
4.6.3.2 The Extended Specification 107
4.7 Conclusion 111
Appendix A 112
References 115

Chapter 5 An Examination of the Efficiency of Emerging Options Markets: The Case of the Athens Derivatives Exchange 121
5.1 Introduction 121
5.1.1 Literature Review 121
5.1.2 Scope of Study 125
5.2 Index Returns 126
5.3 Data 129
5.4 Implied Moments 132
5.4.1 Theoretical Framework 132
5.4.2 Empirical Estimation 134
5.4.3 Alternative Specification 139
5.5 Option Returns 141
5.5.1 Call Options 141
5.5.2 Put Options 146
5.5.3 Option Returns Using the Last Trade Price 149
5.6 Estimation of Options' Risk-Adjusted Returns 150
5.6.1 Single-Factor CAPM 150

5.6.2 Extended CAPM 154
5.7 Straddles 157
5.7.1 Delta-Neutral Straddles 157
5.7.2 Delta and Vega Neutral Straddles 159
5.7.3 Risk-Reversals 162
5.8 Comparison with Empirical Findings from Developed Markets 165
5.8.1 Naked Options Positions 166
5.8.2 Option Strategies 168
5.9 Conclusion 169
Appendix B 171
References 175

Chapter 6 Summary and Conclusions **179**

References **183**

Frequently Used Notation

β_c	beta of a call option
β_p	beta of a put option
Δ_c	call delta
Δ_i	delta of option i
Δ_p	put delta
μ_t	mean of the returns' distribution at time t
$\rho_{ind,t}$	average realized index correlation at time t
σ	volatility of returns
σ_{BS}	Black and Scholes ATM implied volatility
C_{BS}	Black and Scholes theoretical price of a call option
C_{CS}	Corrado and Su theoretical price of a call option
cov_t	covariance matrix at time t
$E_t(\cdot)$	expectation operator conditional on information at time t
$f_t^{t+\tau}$	(logarithmic) forward exchange rate at time t for delivery at time $t+\tau$
h_t	conditional variance at time t
i_t	domestic risk-free interest rate
i_t^*	foreign risk-free interest rate
K	option's exercise price
$N(\cdot)$	standard normal cumulative distribution function
p_t	logarithmic price level at time t
q	constant dividend yield
R_c	daily return of a call option
R_f	daily risk-free rate
$R_{ind,t}$	daily return of the index at time t
R_p	daily return of a put option
R_t	daily log-return
$r_{t+h,t}$	intraday return for the time interval t to $t+h$
rp_t^{re}	risk premium at time t under rational expectations
RV_t	realized intraday index volatility at time t
s_t	(logarithmic) spot exchange rate at time t
t	time
var_t	variance at time t
w_i	weight of asset i

List of Acronyms and Abbreviations

ACSPI	Athens Composite Share Price Index
ADEX	Athens Derivatives Exchange
AMEX	American Stock Exchange
ARCH	Autoregressive Conditional Heteroscedasticity
ATM	at-the-money
CAPM	Capital Asset Pricing Model
CIP	Covered Interest rate Parity
CRSP	Center for Research in Security Prices
DJIA	Dow Jones Industrial Average
EGARCH	Exponential General Autoregressive Conditional Heteroscedasticity
FIGARCH	Fractionally Integrated General Autoregressive Conditional Heteroscedasticity
FM-LAD	Fully Modified Least Absolute Deviations
GARCH	General Autoregressive Conditional Heteroscedasticity
GARCH-M	GARCH in Mean
GJR	Glosten, Jagannathan and Runkle
GPH	Geweke & Porter-Hudak
HAC	Heteroscedasticity and Autocorrelation Consistent
IGARCH	Integrated General Autoregressive Conditional Heteroscedasticity
ITM	in-the-money
JIT	Jensen's Inequality Terms
LAD	Least Absolute Deviations
LIBOR	London Interbank Offered Rate
LR	Likelihood Ratio
LRT	Likelihood Ratio Test
NASDAQ	National Association of Security Dealers Automated Quotation
NYSE	New York Stock Exchange
OLS	Ordinary Least Squares
OTM	out-of-the-money
RND	Risk Neutral Density
S&P	Standard & Poor's
SSE	Sum of Squared Errors
TAQ	Trade And Quotation
UIP	Uncovered Interest rate Parity
WRDS	Wharton Research Data Services

List of Figures

Figure 2.1	DJIA Spot Level
Figure 2.2	DJIA Average Realized Correlation
Figure 2.3	DJIA Realized Volatility
Figure 2.4	Correlogram of DJIA Average Realized Correlation
Figure 2.5	Correlogram of DJIA Realized Volatility
Figure 2.6	Scatterplot of Correlation Changes on Volatility Changes
Figure 4.1	Spot Exchange Rate
Figure 4.2	Forward Exchange Rate
Figure 4.3	Correlogram of Exchange Rate
Figure 4.4	Correlogram of Exchange Rate Returns
Figure 4.5	Nominal Interest Rate Differential
Figure 4.6	Forward Premium
Figure 4.7	Spot Rate's Implied Volatility
Figure 4.8	Rolling OLS Estimations under the Standard Specification
Figure 4.9	Rolling OLS Estimations under the Extended Specification
Figure 4.10	Rolling FM-LAD Estimations under the Standard Specification
Figure 4.11	Rolling FM-LAD Estimations under the Extended Specification
Figure 5.1	FTSE/ASE-20 Index
Figure 5.2	FTSE/ASE-20 Daily Returns
Figure 5.3	FTSE/ASE-20 Daily Returns Distribution
Figure 5.4	FTSE/ASE-320 beta over time
Figure 5.5	Option Observations per Trading Day
Figure 5.6	Implied Volatility (across maturities)
Figure 5.7	Implied Skewness (across maturities)
Figure 5.8	Implied Kurtosis (across maturities)
Figure 5.9	Standardized Implied Volatility
Figure 5.10	Standardized Implied Skewness
Figure 5.11	Standardized Implied Kurtosis
Figure 5.12	Implied Moments from Alternative Methodology
Figure 5.13	Sum of Squared Errors

List of Tables

Table 2.1	List of DJIA Constituent Stocks
Table 2.2	Summary Statistics of the Daily Returns on the DJIA and its Constituents
Table 2.3	Asymmetric Realized Volatility Regressions for the DJIA and its Constituents
Table 2.4	Estimates of Volatility Parameters for the DJIA and its Constituents using the standard GJR model
Table 2.5	Summary Statistics for the DJIA Average Realized Correlation and Realized Volatility
Table 2.6	Intertemporal Relationship between $\Delta\rho_{ind}$, ΔRV and DJIA Log-Returns
Table 2.7	Granger Causality Tests for the Daily Changes in ρ_{ind} and RV
Table 2.8	Asymmetric Volatility Regressions Controlling for Correlation Dynamics
Table 2.9	Estimates of Volatility Parameters for the DJIA using the extended GJR-C model
Table 3.1	Composition of the DJIA
Table 3.2	Constituents Capitalisation
Table 3.3	Summary Statistics of the Daily Returns on the DJIA and its Constituents
Table 3.4	Estimation the Standard GJR
Table 3.5	Estimation of the GJR-I
Table 4.1	Stationarity Tests
Table 4.2	OLS Results
Table 4.3	FM-LAD Estimator Results
Table 5.1	Option Observations per Calendar Day
Table 5.2	Implied Moments
Table 5.3	Strike Groups
Table 5.4	Summary Statistics for Call Options and their Daily Returns (closing price)
Table 5.5	Summary Statistics for Put Options and their Daily Returns (closing price)
Table 5.6	Summary Statistics for Options and their Daily Returns (last trade price)
Table 5.7	Estimated Regression Coefficients of the CAPM
Table 5.8	Estimated Regression Coefficients of the extended CAPM
Table 5.9	Daily Returns of Delta Neutral Straddles
Table 5.10	Daily Returns of Delta and Vega Neutral Straddles
Table 5.11	Risk-Reversals

Chapter 1
Introduction

The accurate measurement and forecasting of volatility have long been of great interest to academic researchers as well as to finance practitioners. The importance of volatility in finance stems mainly from the traditional principle of agents allocating resources into alternative investments based on the risk-return relationship. A rational investor combines securities into a portfolio under the dual criterion of maximizing the expected return and of minimizing the risk of the position. Although the rationality assumption has been questioned by a relatively recent field in the finance literature, usually referred to as 'behavioural finance', the proposition that investors are risk-averse and, as such, interested in the risk of an investment has rarely been challenged.

Given the central role that risk plays in the financial decision making process, it is hardly surprising that a wide body of theoretical and empirical studies have focused on volatility measurement and prediction. The concept of risk is directly linked to volatility since the most common way of establishing an investment's risk is by estimating the volatility of the investment's returns. High volatility is, then, associated with a high level of risk in the sense that the probability of large profits or losses by a specific security is more pronounced compared to low volatility, where the security's return is likely to deviate only moderately from its expected value.

With few exceptions, empirical studies until the 1990s used to proxy volatility as the standard deviation (or variance) of daily returns. Sampling asset returns at a daily basis, although a straightforward task with reasonable at the time data requirements, suffered from certain limitations. For instance, the variance of daily asset returns is considered to provide rather noisy estimates of return volatility and, more importantly, it ignores important information in the intraday path of stock prices. The possibility to address the above concerns arose when high-frequency stock data became more widely available and affordable. Since this development, a significant amount of research has been conducted on high-frequency estimates of volatility, proxied by the sum of intraday squared returns, with ample evidence of high-frequency estimates being superior to low-frequency ones in measuring and forecasting volatility. In addition to complications arising from the size of the high-frequency datasets, though, such as an increased demand for computing power and more elaborate programming, the qualitative characteristics of this data posed new challenges to researchers. More specifically, the bid-ask bounce and other microstructure effects as well as intraday patterns have to be examined and accounted for in order to extract an unbiased proxy of volatility. Despite these complications, the use of high-frequency data in volatility analysis has become a widely adopted technique, with direct implications for risk management, portfolio selection and derivatives trading.

In terms of volatility forecasting, two distinct frameworks have been developed in the literature, which express significantly different lines of thought. The first framework describes the volatility process of a security by fitting an historical model to the security's past returns. This class of historical models are generally termed *Autoregressive Conditional Heteroscedasticity* (ARCH) models, although a variety of specifications has been proposed based on different assumptions about the returns generating process and, thus, with different acronyms.

In contrast, the second framework uses forward-looking information embedded in option prices to extract an estimate of the market's expectation of future volatility. This alternative method of predicting the volatility of asset returns usually relies on the assumptions of a specific option pricing model and, given that implied estimates capture the aggregate investors' expectation of the underlying's future distribution, it theoretically incorporates all relevant information including that contained in historical returns. However, the literature has yet to reach a consensus as to which framework is superior in forecasting volatility, as demonstrated by the subsequent significant interest and the body of research in both fields.

The emergence of ARCH models attempted to accommodate the widespread empirical findings of volatility being serially correlated and moving through time. This directly challenged the validity of the 'constant volatility' assumption, suggesting that changes in volatility might be predictable due to a nonlinear dependence and not necessarily affected by changes in exogenous variables. ARCH specifications model the conditional variance of returns, as opposed to the unconditional variance, as being dependant on past errors, i.e. past returns, and characterized by clusters of high and low observations. In other words, volatility within the ARCH framework responds to innovations in the asset's returns process while periods of high (low) volatility are more likely to be followed by high (low) volatilities.

The simplest specification from the universe of ARCH models is the ARCH(1) which was presented in Engle (1982). The return distribution at time t in Engle's model is conditional on all previous returns and normal with a constant mean μ, while the time-varying conditional variance h_t at time t depends only on the previous return R_{t-1}. Arguably the best known model is the Generalized ARCH (GARCH) which was independently developed by Bollerslev (1986) and Taylor (1986), and in which the conditional variance at t is also dependant on its lagged values in addition to past returns. Although in its simplest and most commonly used form the GARCH(1,1) specification describes volatility as a function of past returns R_{t-1} and of the past value of the conditional variance h_{t-1} at the first lag, the number of lags included for these variables has varied across different studies and ultimately depends on the time-series under examination. The popularity of the simple GARCH(1,1) specification in empirical research is attributed to the fact that, although it only has four parameters and is, therefore, easily estimated, it has been generally found to have comparable accuracy to that of more complex models. Moreover, its results are typically consistent with most stylized facts, particularly with respect to daily returns.

A potential limitation of the GARCH(1,1) model, however, stems from its inability to accommodate the widely observed asymmetric response of volatility to past returns of opposite signs. This conditional correlation of volatility changes to price changes has been termed the 'volatility asymmetry' phenomenon and it describes the empirical finding of volatility increases following negative returns being higher on average than volatility decreases after positive returns of similar magnitude. Volatility asymmetry, which is particularly pronounced in equity markets, was incorporated into a specification presented in Glosten, Jagannathan and Runkle (1993) by placing different weights in positive and in negative residuals. This modified GARCH, referred to as the GJR model, separates the effects of past returns of opposite signs through the use of a dummy variable for lagged negative returns, and subsequent empirical research has shown the GJR to provide a good fit to daily equity returns and to capture the asymmetric impact of signed price changes on the volatility process.

An alternative and equally popular asymmetric GARCH specification is the Exponential GARCH (EGARCH) by Nelson (1991). Similarly to the GJR, the EGARCH model is asymmetric since it contains two parameters for lagged innovations to measure the 'size effect' and the 'sign effect'. However, contrary to the GJR as well as to other GARCH models, the EGARCH relaxes the non-negativity restrictions on the coefficients and the conditional variance is expressed as a multiplicative function of lagged returns rather than as an additive function. Finally, the significant interest that historical volatility models have attracted in the literature has led to the development of a wide class of more complicated and specialized GARCH models that are based on particular properties of the returns generating process. Some of the most commonly used specifications include the Integrated GARCH (IGARCH) where the persistent parameters sum up to unity, the GARCH-in-mean (GARCH-M) where heteroscedasticity is introduced into the mean equation as well, and the Fractionally Integrated GARCH (FIGARCH) which implies long memory for volatility shocks.

The second line of research, as was previously mentioned, examines contingent claims in order to extract information about the future distribution of the underlying assets. Option prices, in particular, are considered to directly reflect the market's expectation of the underlying's distribution for different time horizons and are, thus, a forward-looking measure of investors' beliefs about the future realization of asset prices. Furthermore, the fact that historical prices are readily observable and available to market participants when options are quoted suggests that option prices incorporate all relevant information, including that of historical volatility models. Consequently, option prices should, in principle, contain incremental information about the underlying, in addition to that embedded in historical models, and option-implied estimates should reasonably be expected to provide more accurate forecasts of future volatility compared to those obtained by historical models.

The informational content of option prices has been heavily researched with respect to extracting implied volatility estimates of the underlying upon the options' expiration. One of the earliest and most common methodologies for inferring the market's expectation of future volatility from options is based on the Black and Scholes (1973) volatility of an at-the-money (ATM) option. Although the use of this single option is partly justified by the fact that trading volume is usually the highest for ATM contracts, the analysis has also been extended to applying the Black and Scholes methodology on the entire set of options available, documenting the commonly observed volatility 'smile' and 'smirk', in currency and in equity index options, respectively (see for instance Rubinstein (1985), Taylor and Xu (1994) and Dumas, Fleming and Whaley (1998)).

The above implied volatility measures are subject to the assumptions of the Black and Scholes option pricing formula, the most restrictive of which is arguably that of lognormality in the underlying's return distribution. Given that equity returns are typically found to deviate from lognormality, subsequent research has attempted to create a more flexible framework of extracting option implied information. The Corrado and Su (1996) option pricing formula constitutes such an example, where the normality assumption is relaxed to allow for non-normal third and fourth moments in the underlying's returns. The Corrado and Su framework is practically a modification of the standard Black and Scholes formula that accounts for non-zero skewness and excess kurtosis and, in addition to being more consistent with the stylized facts on equity returns, it also has the advantage of simultaneously estimating the next two implied moments of the returns' distribution as opposed to inferring only implied volatility.

Although empirical studies until the late 1990s were based on specific option pricing models to extract implied volatilities from the prices of options, the related literature has recently moved towards model-free estimates. When a specific model, such as the Black and Scholes or the Corrado and Su model, is used to infer the second moment of the returns' distribution, one jointly assumes the validity of the model as well as the efficiency of the options market. This limitation was addressed in Britten-Jones and Neuberger (2000) where it was shown that the future volatility of asset returns can be estimated from option prices without relying on a pricing model. More specifically, Britten-Jones and Neuberger (2000) demonstrated that implied volatility is linked to the expected sum of squared returns under the risk-neutral measure and that it is completely specified by a set of out-of-the-money (OTM) options. This methodology provides a 'model-free' estimate of the underlying's future volatility without making any assumptions about the underlying's returns distribution, and it has become increasingly popular in subsequent empirical papers.

Other studies have expanded the literature by focusing on the entire *Risk Neutral Density* (RND) of the underlying that is specified by a set of option prices, instead of extracting only the second moment of the distribution. Similarly to the implied volatility estimation techniques, certain assumptions about the underlying need to be made in order to infer the RND from option prices. A relatively common way of modelling the returns' generating process is as a mixture of two lognormal distributions, as opposed to the single lognormal of the Black and Scholes formula. The double lognormal has been found to provide a close fit to market prices for a variety of indices and stocks, although alternative smoothing methodologies, such as cubic splines and implied binomial trees, have also been shown to provide similarly close fits to market data.

Finally, admittedly the vast majority of papers on option-implied information have focused on extracting the first moments or the entire RND of a univariate distribution. However, the 2000s have experienced an increased interest in the informational content of 'implied association'. The concept of implied association refers to the dependence pattern that ensures the price consistency between multivariate and univariate options, and most of the related studies examine it from the perspective of 'implied correlation', i.e. a single unconditional estimate of implied dependence per time-period. Although option-implied correlation had been incorporated in empirical papers during the early 2000's, a significant contribution was made by a recent paper by Driessen, Maenhout and Vilkov (2008) that presented a straightforward framework for inferring an estimate of the average implied correlation among an index's components using the prices of options written on the index and on the constituent stocks.

This thesis consists of four empirical studies that are related to the previously mentioned frameworks of modelling and forecasting volatility, i.e. historical GARCH specifications and option-implied estimates. Although these studies fall under the above general line of thought, the thesis was elaborated in a way that each study can be read independently. In this sense, Chapters 2 to 5 constitute discrete exercises that examine different research questions and lead to independent conclusions.

The second Chapter presents and tests an alternative hypothesis for the 'asymmetric volatility' phenomenon in index returns. The two most common explanations of volatility asymmetry, namely the 'leverage effect' and 'volatility feedback', fail to accommodate the widely reported finding of index volatility being significantly more asymmetric than the volatilities of individual stocks. Motivated by this discrepancy, the second Chapter explores the possibility of changes in the average realized correlation among the index's constituents driving the asymmetric response of the parent index's volatility to past returns of opposite signs. This alternative explanation is directly linked to Rubinstein's (2000) 'diversification effect' which states that an increase (decrease) in the index's average correlation reduces (increases) the benefits of diversification and, therefore, increases (decreases) the index's risk as reflected by a higher (lower) volatility of returns. Given that the index's variance is essentially the weighted sum of individual stock variances plus the cross-correlation terms and that index volatility asymmetry does not appear to stem from asymmetric individual stock volatilities, the dynamics of the average index correlation could potentially reconcile the difference between the volatility processes of these two asset classes. Focusing on the Dow Jones Industrial Average across a ten-year sample period, the second Chapter tests the 'diversification effect' hypothesis by examining the co-movement of the index's average realized correlation with index returns and with index volatility, as well as by estimating an extended GJR specification that includes conditional changes in correlation as an exogenous regressor in the index's conditional variance equation.

The third Chapter shifts the focus from the index's volatility to that of individual stocks, with the emphasis remaining on the 'volatility asymmetry' phenomenon. More specifically, the common empirical finding of individual stock volatilities being less asymmetric than those of indices is analyzed from the perspective of the 'down-market effect' which treats volatility asymmetry as a result of market-level influences rather than of asset-specific innovations. Within this context, the hypothesis is tested that the conditional variances of equities (including equity indices) respond asymmetrically to market 'news' rather than to idiosyncratic 'news', a hypothesis which is consistent with the stylized fact of volatility asymmetry being more pronounced in the case of market indices compared to individual stocks. In order to examine the effect of systematic 'news' on the volatility processes of individual stocks, a modification of the GJR model is estimated across the thirty components of the Dow Jones where lagged signed market returns have replaced firm-specific returns, as well as an extension where the conditional variance responds both to idiosyncratic and to systematic innovations.

The last two Chapters move from empirical applications on historical volatility models to the informational content of option-implied volatility and to the efficiency of options markets. The objective of the fourth Chapter in particular is to re-examine one of the most heavily cited topics in the finance literature, namely the 'forward premium puzzle', using information implied by currency options. This empirical anomaly refers to the widely reported finding that when currency spot returns are regressed on the forward premium, slope coefficients are produced that are systematically less than unity and, in many cases, less than zero. Negative slopes are, therefore, interpreted as evidence against the forward rate being a conditionally unbiased predictor of future spot rates and as a significant violation of *Uncovered Interest Parity*, making excess returns in the foreign exchange markets appear predictable. The fourth Chapter explores the hypothesis of Jensen's Inequality being related to the magnitude of the observed difference between forward rates and the subsequent realizations of spot rates. In contrast to previous papers, though, the Jensen's Inequality Term is proxied by the option-implied variance of the spot rate and the hypothesis of interest is tested by estimating an extended specification where the spot rate's future variance is included as an additional regressor using the Fully Modified Least Absolute Deviations (FM-LAD) estimator.

The use of options to extract information about the future distribution of asset returns is directly linked to the informational efficiency of the options market. Inferring option-implied estimates of volatility, thus, assumes that options are correctly priced to reflect the aggregate market expectation of future volatility. Although previous studies have examined the efficiency of developed options markets, such as the US and the UK, to report that some mispricing exists that is not arbitraged away due to transaction costs, the issue of mispricing in emerging markets so far has been ignored. The fifth Chapter attempts to address this concern by focusing on the developing Greek options market of the Athens Derivatives Exchange (ADEX). The efficiency of the market is examined from the perspective of option returns that are commensurate with the underlying risks and consistent with theoretical pricing models, so that the market is considered to be efficient if options are priced to compensate investors for risk and do not offer abnormal returns. The hypothesis of interest is that, given the globalized marketplace and the fact that most of the trading volume in Greece is attributed to specialized international investors, the emerging market of ADEX is comparably efficient with respect to developed markets. The alternative hypothesis is partly motivated by the findings of Santa-Clara and Saretto (2009) and states that higher transaction costs combined with thinner trading in ADEX will widen the no-arbitrage bands and are likely to be associated with more pronounced mispricing. The efficiency of the Greek options market is examined by two commonly used measures: the significance of CAPM alphas that effectively measure the options' risk-adjusted returns, and the returns of positions that are immune to the underlying risks (i.e. delta and delta-vega neutral straddles). The results are then contrasted to those typically obtained in the case of the US and the UK developed options markets to determine the comparative efficiency of the ADEX.

Finally, Chapter 6 provides an overview of the main findings and contributions of the present thesis. The results of each Chapter are summarized separately following a general conclusion regarding the thesis' unifying topic.

Chapter 2

The Role of Realized Correlation Dynamics in Explaining Volatility Asymmetry in Dow Jones Index Returns

2.1 Introduction

2.1.1 Literature Review

It has become common practice in the finance literature to model stock return volatility as negatively correlated with stock returns. In addition to this negative correlation, it is widely accepted that, dependent on the asset of interest, the above relationship is asymmetric and conditional on the return's sign. The above 'asymmetric volatility' phenomenon has been commonly referred to as the 'leverage effect', due to the fact that a large part of the early studies that examined this relationship attributed volatility changes to changes in a firm's leverage. One of the earliest papers to discuss this phenomenon is Black (1976) who demonstrates that volatility responses to negative shocks are fundamentally different than those to positive shocks. More specifically, asset volatility following a negative return is systematically found to be higher compared to volatility that follows a positive return of similar magnitude. Black suggests that this asymmetry is potentially attributed to the fact that negative returns reduce the value of a firm's equity relative to its debt, thereby making the overall firm riskier, a fact that is reflected in higher levels of stock price volatility.

Black's hypothesis of leverage increases turning into increases in asset volatility has been empirically tested by Christie (1982) who finds that, although changes in leverage are significantly related to changes in volatility, the debt-to-equity ratio fails to fully explain the observed effect. Similar conclusions are reached by Figlewski and Wang (2000) who examine a large sample of stocks included in the Standard and Poor's 100 Index, as well as the index itself. Using direct measures of a firm's debt-to-equity ratio, Figlewski and Wang (2000) show that, in contrast to theoretical predictions, changes in leverage have different impacts on realized and on implied volatility changes within their sample. Moreover, changes in leverage alone fail to fully explain volatility changes, especially in the case of up markets, with the authors concluding that the 'leverage effect' is mostly not attributed to leverage but '...should more properly be termed as a down market effect'.

Campbell and Hentschel (1992) and French, Schwert and Stambaugh (1987) propose volatility feedback as an alternative explanation for volatility asymmetry. When volatility increases at time *t*, expected risk premia at $t+1$ also increase to compensate investors for bearing more risk. An increase in risk premia, i.e. a higher return at $t+1$, is associated with a decrease in the stock's price at *t*, leading to a negative contemporaneous relationship between returns and unexpected changes in volatility. A recent study by Smith (2007) explores the effect of 'volatility feedback' within the context of the risk-return relationship. The author examines a large sample of US stocks and indices, and finds that including volatility feedback in a stochastic volatility model explains a statistically and economically significant part of the variance of daily returns, while the model's negative conditional correlation between stock returns and volatility closely approximates the actually observed correlation coefficient.

Bekaert and Wu (2000) compare the magnitude of the 'leverage effect' and of 'volatility feedback' in explaining volatility asymmetry and suggest that the latter effect dominates the former. Moreover, Wu (2001) extends the model of Campbell and Hentschel (1992) to incorporate both the leverage effect and the volatility feedback effect as potential explanations, and finds that both effects are determinants of asymmetric volatility, with volatility feedback being statistically and economically significant.

Bollerslev, Litvinova and Tauchen (2006) further examine asymmetric conditional variance using high-frequency data, with particular focus on the causality relationship between volatility and returns. More specifically, the 'leverage effect' explanation suggests that stock returns cause changes in volatility through changes in the debt-to-equity ratio and, consequently, in the underlying risk of the asset. On the other hand, the 'volatility feedback' explanation reverses this relationship, with unanticipated volatility changes causing changes in the current required return of the asset to compensate for the new level of underlying risk. Bollerslev *et al* (2006) use intraday returns of futures written on the S&P 500 from January 1988 to March 1999, and report a strong negative correlation between volatility and current and lagged index returns, while the correlation between returns and lagged volatility is found to be very close to zero. Their results are interpreted as evidence of a prolonged leverage effect in intraday returns that lasts for several days, combined with '...an almost instantaneous volatility feedback effect'.

The GJR model, developed by Glosten, Jagannathan and Runkle (1993), is one of the most commonly used to account for the asymmetric volatility phenomenon. The GJR model specifically allows for an asymmetric response of volatility to lagged returns, with the conditional variance depending on the sign as well as on the magnitude of past returns. Another popular asymmetric model is Nelson's (1991) Exponential GARCH (EGARCH), which was modified by Cheung and Ng (1992) to include a lagged price variable in the conditional variance. The subsequent empirical testing of these asymmetric models has suggested that volatility dynamics are significantly different for the parent index compared to its constituent stocks. For instance, Lamoureux and Lastrapes (1990) investigate the effect that structural changes have on the volatility processes of thirty randomly selected US stocks as well as on that of the CRSP value-weighted index, and find that the GARCH model appears to overstate the persistence in variance by failing to take into account structural shifts in the model. When dummy variables are introduced that permit a non-stationary unconditional variance, the persistence estimates are lower compared to those of a simple GARCH(1,1). However, this decline in variance persistence is relatively greater for individual stocks (average persistence falls from 0.97 to 0.82) than for the index (from 0.99 to 0.96), indicating that volatility dynamics are significantly different for these two asset categories.

Kim and Kon (1994) provide additional evidence for different asymmetric responses of the conditional variance to returns in the case of indices and of individual stocks. By examining a sample of individual stocks (the thirty constituents of the Dow Jones Industrial Average) and three indices (S&P 500, CRSP Equally Weighted, and CRSP Value Weighted) from 1962 to 1990, they find that volatility asymmetry is more pronounced for index returns compared to returns of individual stocks. This finding is also supported by Tauchen, Zhang and Liu (1996) who compare the volatility processes of four large-capitalization stocks to those traditionally observed for aggregate market indices, and conclude that the 'leverage effect' is present but relatively weaker in the conditional variance of individual stocks.

Blair, Poon and Taylor (2002) examine the volatility processes of the S&P 100 and of its constituents from 1983 to 1992, with particular focus on the asymmetry with respect to past returns as well as on the effects of the October 1987 crash. Using an extension of the GJR model, they report that the volatility processes of the index and of the majority of stocks indeed support the existence of the so-called 'leverage effect'. More importantly, though, the asymmetric response of the conditional variance to lagged returns appears to be significantly more pronounced for the index compared to individual stocks.

Stivers, Dennis and Mayhew (2006) provide further evidence on the differences between the volatility processes of indices and of individual stocks by separating the effect of changes in idiosyncratic and in systematic volatility. Their study relies on the key assumption that implied volatility represents an observable proxy for the expected return volatility and, thus, defines the 'asymmetric volatility' phenomenon as the relationship between returns and innovations in implied volatility. Examining a sample of daily returns for the S&P 100 index and for fifty large US firms across the period 1988 to 1995, the authors find that the daily returns of individual stocks are weakly correlated with changes in idiosyncratic volatility, proxied by firm-specific implied volatility. This finding, combined with the fact that index returns exhibit a strong negative comovement with changes in systematic volatility, as measured by the index's implied volatility, is interpreted as evidence of a more pronounced 'asymmetric volatility' effect for the index compared to individual stocks. Moreover, firm-specific stock returns are found to be significantly correlated with changes in the index's implied volatility, i.e. with innovations in systematic volatility, suggesting that the 'asymmetric volatility' phenomenon describes the effect of market-level influences rather than a 'leverage effect' or 'volatility feedback', consistent with the results by Figlewski and Wang (2000).

Given that the index's variance comprises of the constituents' variances and of pairwise correlations among the individual stocks, and that asymmetric index volatility cannot be fully attributed to asymmetries in the volatility processes of the constituents, it has been suggested that volatility asymmetry in the index might be caused by changes in correlations. The above hypothesis is directly related to Rubinstein's (2000) 'diversification effect' which describes the fact that cross-correlations among individual securities tend to increase when the index goes down, reducing the benefits of diversification when it is needed the most. Within this framework, a negative index return is associated with an increase in the average correlation of its constituents and leads to higher index volatility. Therefore, Rubinstein's 'diversification effect' provides a theoretical justification for a positive relationship between changes in average correlation and index volatility, which is the focus of the present study.

In fact, Cappiello, Engle and Sheppard (2006) document a similar negative relationship between index returns and correlations for a sample of FTSE All-World Indices (21 countries) and 5-year average maturity bond indices (13 countries). Cappiello *et al* find that correlations in equity and bond indices exhibit asymmetric responses to bad news relative to good ones, although equities appear to be more significantly affected than bonds. Furthermore, Kroner and Ng (1998) report an asymmetric response of covariance between large and small firms, with the conditional covariance tending to be higher following bad news about large firms compared to good news. Finally, Driessen, Maenhout and Vilkov (2008) use a measure of the average implied correlation among an index's constituents to examine the commonly reported finding of index options bearing a larger (negative) risk premium compared to options written on individual stocks. Using options data on the S&P 100 and on its components from January 1996 to December 2003, Driessen, Maenhout and Vilkov (2008) show that the risk stemming from changes in the average implied correlation of the index is priced in the US market, concluding that correlation risk '... contributes to the variance risk premium of the index, but is not present in the individual stock options in its orthogonal part'.

2.1.2 Scope of Study

The aim of this Chapter is to investigate whether changes in correlations among the constituent stocks can explain the volatility asymmetry of the parent index. Despite the large amount of empirical evidence documenting an asymmetric response of the conditional variance to lagged returns of different signs in the case of indices, and to a relatively smaller degree in individual stocks, the literature has still to reach a consensus on the source of this phenomenon. The intuition behind changes in correlations driving the 'leverage effect' stems from the fact that the index's variance is defined as the sum of the variances of its components plus the correlation terms among individual stocks. Since volatility asymmetry is significantly more pronounced in the index compared to the constituent stocks, it could be the case that the asymmetric response of index volatility to lagged returns is driven, at least partly, by changes in correlations.

Rubinstein's (2000) 'diversification effect' suggests that correlations among stocks increase subsequent to a negative index return, indicating a positive relationship between changes in correlations and index volatility. Blair, Poon and Taylor (2002) also highlight significant differences in the volatility processes of an index and of its components. However, to the best of my knowledge, the extent to which correlation dynamics can explain the 'asymmetric volatility' phenomenon has not been explicitly tested. The present study attempts to fill this gap by, among other empirical tests, estimating an extension of the GJR asymmetric model on DJIA returns that specifically accounts for changes in the average correlation among the index's constituents.

Overall, a set of initial results is presented that supports the 'diversification effect' hypothesis, with changes in the average DJIA realized correlation being negatively related to index returns. This relationship is found to be asymmetric, since correlation increases by more following a negative return compared to its decrease following a positive return of similar magnitude. Some additional evidence for the diversification hypothesis is provided by the positive relationship between the average realized correlation and the index's realized intraday volatility, which suggests that increases (decreases) in the correlations among the index's constituents result in increases (decreases) of the index's level of riskiness. Finally, when an extension of the GJR specification is estimated for DJIA returns, the asymmetry in the conditional variance's responses to past returns of different signs is reduced, albeit still present, and the coefficient for correlation changes is highly significant in 'down' markets while insignificant in 'up' markets. In appears that, although the negative relationship between volatility and the direction of market movements strongly holds, its asymmetric property that relates to the magnitude of index returns is weakened after accounting for the average correlation's dynamics. The above results can be interpreted as initial evidence for conditional innovations in the average correlation among the index's components at least partly driving the observed volatility asymmetry in Dow Jones index returns.

The remaining of the Chapter is organized as follows. Section 2.2 presents the data used in this Chapter. Section 2.3 describes the methodology for estimating the index's realized volatility and the realized average correlation among the index's components from high-frequency data on the individual stocks. Section 2.4 analyzes the asymmetric response of index volatility to lagged index returns of different signs. Section 2.5 discusses the time-series properties of the average realized correlation of the Dow Jones Industrial Average and the relationship between the average realized correlation and the index's realized volatility. Section 2.6 examines the extent to which correlation dynamics can explain volatility asymmetry in DJIA returns. Finally, Section 2.7 concludes.

2.2 Data

The high-frequency data used in this study is from the New York Stock Exchange (NYSE) Trade And Quotation (TAQ) database. The TAQ data files provide continuously recorded time-stamped prices, volumes and bid-ask quotes of particular transactions for stocks traded in NYSE, the American Stock Exchange (AMEX), and the National Association of Security Dealers Automated Quotation system (NASDAQ). The specific sample used in this empirical analysis focuses on the Dow Jones stocks that are traded in NYSE, where trading hours extend from 9:30 EST until 16:05 EST. Relevant dividend distributions are also included in the TAQ dataset.

The sample period starts on 30 March 1998 and ends on 30 March 2007, for a total of 2,265 trading days. Table 2.1 lists all 37 companies that were members of the DJIA index at any stage during the sample period, including company names, ticker symbols and, where applicable, name changes, dates of entry and dates of exit from the index. Between March 1998 and March 2007, the composition of DJIA changed twice, namely on 1 November 1999 and on 8 April 2004, with four replacements taking place on the first date and three on the second. In total, 23 companies remained in the DJIA throughout the entire sample period (Survivors), while 7 companies left (Leavers) and 7 were introduced in the index (Entrants). Finally, one company (AT&T) exited the index, with its ticker re-entering after a merger with another index member (SBC).

Figure 2.1

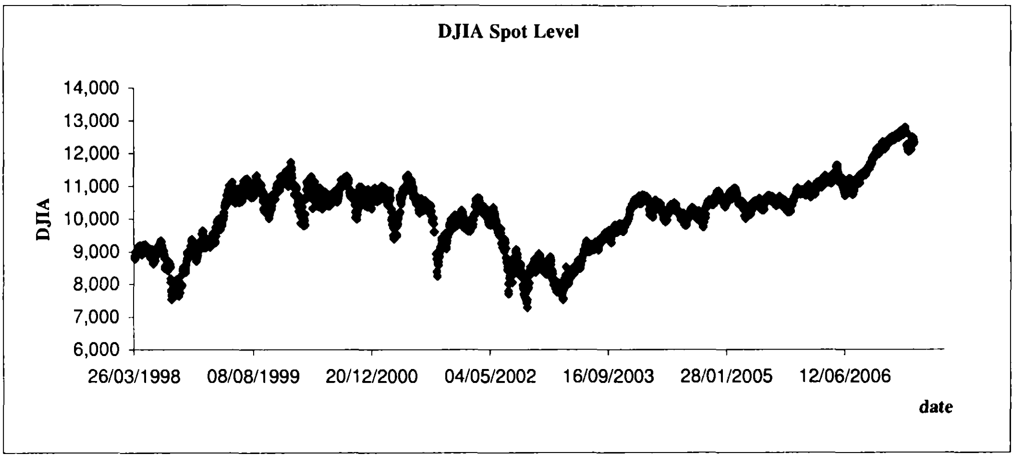

The closing spot levels of the index as well as of its constituent stocks for the sample period were obtained through DataStream. The above prices have been adjusted for dividends and for changes in capital structure, and are used to compute log-returns at a daily frequency.

Table 2.1

List of DJIA Constituent Stocks

Company Name	Ticker	Date of Entry	Date of Exit
AlliedSignal Incorporated (Honeywell International)	ALD (HON)		
Aluminum Company of America (Alcoa)	AA		
American Express Company	AXP		08/04/2004
American International Group Incorporated	AIG	08/04/2004	
AT&T Corporation	T		
Boeing Company	BA		
Caterpillar Incorporated	CAT		
Chevron	CVX		01/11/1999
Coca-Cola Company	KO		
DuPont	DD		
Eastman Kodak Company	EK		08/04/2004
Exxon Corporation (Exxon Mobil)	XOM		
General Electric Company	GE		
General Motors Corporation	GM		
Goodyear	GT		01/11/1999
Hewlett-Packard Company	HPQ		
Home Depot Incorporated	HD	01/11/1999	
Intel Corporation	INTC	01/11/1999	
International Business Machines	IBM		
International Paper Company	IP		08/04/2004
Johnson & Johnson Company	JNJ		
J.P. Morgan & Company	JPM		
McDonald's Corporation	MCD		
Merck & Company Inc	MRK		
Microsoft Corporation	MSFT	01/11/1999	
Minnesota Mining & Mfg (3M Company)	MMM		
Pfizer Incorporated	PFE	08/04/2004	
Philip Morris Companies Inc (Altria Group, Incorporated)	MO		
Procter & Gamble Company	PG		
SBC Communications (AT&T)	SBC (T)	01/11/1999	
Sears & Roebuck Company	S		01/11/1999
Travelers Group (Citigroup)	TRV (C)		
Union Carbide	UK		01/11/1999
United Technologies Corporation	UTX		
Verizon Communications Incorporated	VZ	08/04/2004	
Wal-Mart Stores Incorporated	WMT		
Walt Disney Company	DIS		

This Table provides a list of the companies that were members of the Dow Jones Industrial Average index (DJIA) at any point during the sample period of March 1998 to March 2007. The parentheses in the 'Company Name' field refer to name changes within the sample period.

2.3 Estimation of the Realized Volatility and the Average Realized Correlation of the DJIA

Andersen, Bollerslev, Diebold and Ebens (2001a) provide a detailed discussion on the theoretical justification for using the sum of intraday squared returns as a measure of realized stock volatility (see also Andersen, Bollerslev, Diebold and Labys (2001b)). This Section summarizes the above discussion and describes the methodology for estimating the realized index volatility and the average realized correlation among the index's components using high-frequency stock data. Assume a $Nx1$ vector of the logarithmic price process p_t that follows a multivariate continuous time stochastic volatility diffusion:

$$dp_t = \mu_t d_t + \Omega_t dW_t \tag{2.1}$$

where W_t is a standard N-dimensional Brownian motion, Ω_t is strictly stationary, and the time unit h is normalized to a trading day. The period return $r_{t+h,t}$ is then defined as the logarithmic difference between the spot prices at t and at $t+h$:

$$r_{t+h,t} \equiv p_{t+h} - p_t \tag{2.2}$$

Conditional on the sample path realization of μ_t and Ω_t, the distribution of the continuously compounded returns $r_{t+h,t}$ is then:

$$r_{t+h,t} \mid \sigma\{\mu_{t+\tau}, \Omega_{t+\tau}\}_{\tau=0}^{h} \sim N\left(\int_0^h \mu_{t+\tau} d\tau, \int_0^h \Omega_{t+\tau} d\tau\right) \tag{2.3}$$

where $\sigma\{\mu_{t+\tau}, \Omega_{t+\tau}\}_{\tau=0}^{h}$ refers to the σ-field generated by the sample paths of the drift $\mu_{t+\tau}$ and the diffusion matrix $\Omega_{t+\tau}$ for $0 \leq \tau \leq h$. Hence, the integrated diffusion matrix provides '... a natural measure of the true latent h-period volatility'. Andersen *et al* (2001a) then suggest that, under weak regularity conditions,

$$\sum_{j=1,\ldots,[h/\Delta]} r_{t+j\Delta,\Delta} r'_{t+j\Delta,\Delta} - \int_0^h \Omega_{t+\tau} d\tau \rightarrow 0 \tag{2.4}$$

as $\Delta \rightarrow 0$, i.e. as the sampling frequency of the returns increases. Thus, dependent on the sampling interval, the sum of sufficiently finely sampled intraday returns can provide a direct measure of ex post realized volatility that is asymptotically free of measurement error.

The methodology in Andersen and Bollerslev (1997a) and in Andersen *et al* (2001a) is adopted in order to extract a measure of realized volatilities and pairwise correlations of the DJIA's components using high frequency quotes. Similarly to the above papers, time-series of artificial intraday returns for each stock are constructed at five-minute intervals. The five-minute sampling frequency is typically considered to be short enough so that the summation in (2.4) closely approximates the integrated volatility, and long enough to minimize the noise stemming from market-microstructure effects.

Given that trading hours in NYSE extend from 9:30 EST until 16:05 EST, one trading day can be decomposed in 79 five-minute intervals, such that Δ is equal to 1/79, or roughly 0.0127. Intraday spot levels are measured as the midpoint of the best bid and the best ask prices recorded at or immediately before the 80 five-minute marks (note that best quotes are used instead of prices at which actual trades occurred). The corresponding 79 five-minute returns in a trading day are then computed as the logarithmic differences between consecutive five-minute marked spot prices.

Due to the use of a discrete sampling interval to approximate the continuous volatility process, Andersen *et al* (2001a) suggest that the presence of negative serial correlation in the returns series as well as the inherent bid-ask spread are likely to bias the estimation of the above volatility measure. In order to avoid or at least minimize the dependence in the mean of the observed intraday log-differences, demeaned returns are estimated by fitting a simple MA(1) model for each of the five-minute series across the full ten-year sample. Also note that, similarly to previous empirical studies, high-frequency returns refer only to intraday log-differences in spot prices, i.e. spot changes within the trading day, excluding any overnight price changes. Although this approach admittedly results in some loss of information with respect to the underlying returns series, it is associated with a less noisy and well-behaved time-series. For notational simplicity, the resulting five-minute demeaned return for stock i that is recorded at day t and at the five-minute interval k is denoted by $r^i_{t+k\Delta,\Delta}$. Hence, the realized daily covariance matrix cov_t is given by:

$$\text{cov}_t(i,j) \equiv \sum_{k=1,\ldots,1/\Delta} r^i_{t+k\Delta,\Delta} r^j_{t+k\Delta,\Delta} \tag{2.5}$$

where the elements in the diagonal of $cov_t(i,j)$ refer to the intraday realized variances $v^2_{i,t} \equiv \{\text{cov}_t(i,i)\}$ of the thirty stocks in the DJIA. Similarly, the intraday realized covariances between stocks i and j at time t are given by the elements of cov_t outside the diagonal, with the intraday realized correlations denoted by $corr_t(i,j)$:

$$corr_t(i,j) \underset{i \neq j}{\equiv} \frac{\{\text{cov}_t(i,j)\}}{v_{i,t} v_{j,t}} \tag{2.6}$$

Finally, the weighted average index correlation $\rho_{ind,t}$ at time t refers to the average correlation across all possible pairs of the DJIA components, scaled by each stock's weight in the composition of the index:

$$\rho_{ind,t} \equiv \sum_{i=1,j=1,i \neq j}^{N} w_i w_j corr_t(i,j) \tag{2.7}$$

where N gives the number of stocks included in the index, i.e. thirty component stocks in the case of the Dow Jones Industrial Average, and w_i refers to the weight of stock i in the index's composition. Since the DJIA is a value-weighted index, the weight w_i of stock i is essentially the ratio of the stock's price divided by the sum of the prices of all thirty components.

In addition to computing realized volatilities for the DJIA constituent stocks and the average realized correlation of the parent index, the above time-series of demeaned five-minute returns of the individual stocks also allow the estimation of the index's realized intraday volatility. More specifically, the index's intraday return $r^{ind}_{t+k\Delta,\Delta}$ at day t and at the five-minute interval k can be easily obtained as the weighted sum of the constituents' five-minute returns in (2.8).

$$r^{ind}_{t+k\Delta,\Delta} \equiv \sum_{i=1}^{N} w_i r^{i}_{t+k\Delta,\Delta} \tag{2.8}$$

Similarly to the methodology described above for the individual stocks, the DJIA realized variance $v^2_{ind,t}$ is proxied by the sum of squared intraday index returns, i.e. $v^2_{ind,t} \equiv \sum_{k=1,\ldots,[h/\Delta]} [(r^{ind}_{t+k\Delta,\Delta})^2]$. Finally, the DJIA's realized volatility at time t is denoted by RV_t and is measured as the (annualized) squared root of $v^2_{ind,t}$.

2.4 Volatility Asymmetry in DJIA Index Returns

This Section analyzes the returns distributions of the Dow Jones Industrial Average and of its constituent stocks, with the main focus on determining whether volatility responds asymmetrically to past returns of opposite signs. First, descriptive statistics of the unconditional distribution for the DJIA are reported and contrasted to those for the 37 individual stocks in the sample. Moreover, asymmetric volatility regressions and the standard GJR model are estimated across all assets.

2.4.1 Daily Returns of the DJIA and of its Constituents

Table 2.2 presents descriptive statistics for the log-returns of the DJIA and of its components during the sample period of 30 March 1998 to 30 March 2007. In addition to the first four moments of the distribution, the Table reports minimum and maximum values, the autocorrelation at the first lag and the Box-Ljung statistics for the first fifteen lags. With respect to individual stocks, the median, lower quartile (L_q) and upper quartile

(U_q) of each statistic are provided, while the last column gives the number of constituents having a statistic greater than that of the index.

As can be seen from the Table, the DJIA mean daily return is 2.12 basis points, compared to 3.51 bps for the median stock. Given the benefits of diversification, it is not surprising that the index's variance (1.22 bps) is significantly lower than that of the median stock (4.50 bps) or even of the first quartile stock (3.42 bps). Furthermore, the DJIA has lower skewness and kurtosis compared to 23 out of a total of 37 stocks, with the vast majority of stocks exhibiting lower minimum values and higher maximums compared to the index. Finally, the index has a small but insignificant negative autocorrelation at the first lag (autocorrelation is -0.01) with most of the stocks also having insignificant first-order autocorrelations ($\rho_{r,1}$ for the median stock is 0.00).

Table 2.2

Summary Statistics of the Daily Returns on the DJIA and its Constituents

Summary Statistics	Index	Stocks			
		Median	L_q	U_q	Greater than Index
Mean x 10^4	2.12	3.51	1.35	5.09	23
Variance x 10^4	1.22	4.50	3.42	5.32	36
Skewness	-0.03	0.08	-0.33	0.23	23
Kurtosis	3.71	5.06	3.11	10.74	23
Max	0.06	0.11	0.10	0.14	36
Min	-0.07	-0.14	-0.18	-0.10	3
$\rho_{r,1}$	-0.01	0.00	-0.02	0.02	24
Q_1	25.93*	22.86	17.68	29.93*	15

This table tabulates summary statistics of the daily returns on the DJIA index and its constituents. The sample period is 30 March 1998 to 30 March 2007. The index's constituents include firms that remained in the DJIA throughout the entire period (Survivors), as well as Leavers and Entrants, for a total of 37 individual stocks. $\rho_{r,1}$ refers to the first order autocorrelation in daily returns. Q_1 is the Box-Ljung statistic (first 15 lags) and * denotes a significant statistic at the 5% confidence level.

This study places particular emphasis on the asymmetric relationship between return volatility and past returns of opposite signs. Panel A of Table 2.3 reports the results of the Engle and Ng (1993) test for negative bias. Similarly to the previous Table, the relevant statistics for the index are reported as well as for the median, lower and upper quartile stocks. The Engle and Ng (1993) methodology tests whether volatility responds asymmetrically to past returns of different signs by regressing the following equation:

$$R_t^2 = \alpha + \beta s_{t-1} R_{t-1} + \varepsilon_t \tag{2.9}$$

where R_t is the asset's daily return at time t, and s_{t-1} is a dummy variable that takes the value of one if the lagged return R_{t-1} is negative and the value of zero otherwise. [1] The Negative Bias statistic is then equal to the t-ratio of the slope coefficient β in the above regression, with a large negative t-ratio suggesting a more pronounced asymmetry of volatility with respect to lagged negative and positive returns. The results presented in Table 2.3 indicate that the degree to which volatility is more sensitive to negative returns compared to positive ones, is significantly more pronounced for the index than for its components. More specifically, the t-ratio of the regression in (2.9) in the case of the DJIA is -6.69, while the median (lower quartile) stock has a slope t-ratio of -5.14 (-6.48). More importantly, 30 out of 37 stocks have statistics that are lower (in absolute terms) than the index's, confirming previous empirical findings of volatility asymmetry being significantly more evident for indices compared to individual stocks.

[1] The upper case variable R_t refers to daily returns using closing prices as opposed to intraday returns which are denoted by the lower case variable $r_{t+h,t}$.

Table 2.3

Asymmetric Realized Volatility Regressions for the DJIA and its Constituents

Panel A: Negative Bias Test

$$R_t^2 = \alpha + \beta s_{t-1} R_{t-1} + \varepsilon_t$$

Index		-6.69
Stocks	Median	-5.14
	L_q	-6.48
	U_q	-3.51
	Greater than Index	30

Panel B: $\Delta RV_t = \alpha + \sum_{i=1}^{5} \beta_i \Delta RV_{t-i} + \gamma^- s^- R_{t-1} + \gamma^+ s^+ R_{t-1} + \varepsilon_t$

	Index	Stocks		
		Median	L_q	U_q
γ^-	-0.868	-0.2444	-0.5189	0.0152
	(-6.94)			
γ^+	-0.668	-0.5315	-0.7246	-0.2479
	(-5.44)			

Panel C: $\Delta RV_t = \alpha + \sum_{i=1}^{5} \beta_i \Delta RV_{t-i} + \gamma_0 R_t + \gamma_{-1} R_{t-1} + \gamma_{|0|} |R_t| + \gamma_{|-1|} |R_{t-1}| + \varepsilon_t$

	Index	Stocks		
		Median	L_q	U_q
γ_0	-0.7826	-0.3801	-0.5403	-0.0593
	(-11.71)			
γ_{-1}	-0.7463	-0.2358	-0.5193	-0.1486
	(-10.84)			
$\gamma_{\|0\|}$	0.6431	3.4092	2.3538	4.2394
	(6.57)			
$\gamma_{\|-1\|}$	0.0230	-0.7364	-1.2614	-0.5193
	(0.23)			

This Table tabulates the results of regressions examining the asymmetric responses of realized volatility to past returns of opposite signs for the DJIA and its 37 constituent stocks. Panel A presents the results of the Engle and Ng (1993) test for negative bias in the returns processes. Panel B presents the results of regressing changes in realized volatility on lagged signed returns. Panel C presents the results of regressing changes in realized volatility on contemporaneous and lagged returns and absolute returns. The sample period is March 1998 to March 2007.

2.4.2 Volatility Asymmetry in Index and in Individual Stocks' Returns

Another methodology that has been frequently employed in the 'leverage effect' literature to detect asymmetries in the volatility process involves a regression of volatility changes on lagged returns. Equation (2.10) describes such as specification, where changes in realized volatility are regressed on signed past returns. RV_t is the realized volatility at time t, ΔRV_t refers to changes in realized volatility at t and is computed as the first difference $(RV_t - RV_{t-1})$, s^- is a dummy variable that takes the value of one if the lagged return R_{t-1} is negative and of zero otherwise, while s^+ is a dummy that takes the value of one if R_{t-1} is positive and of zero otherwise.

$$\Delta RV_t = \alpha + \sum_{i=1}^{5} \beta_i \Delta RV_{t-i} + \gamma^- s^- R_{t-1} + \gamma^+ s^+ R_{t-1} + \varepsilon_t \qquad (2.10)$$

The specification in (2.10) is examined at a daily frequency using the time-series of realized volatility estimated in Section 2.3 from high-frequency data. The first five lags of realized volatility have been included in the regression in an attempt to control for any autocorrelations in the time-series. Within this framework, equation (2.10) distinguishes between the effects of past negative and positive returns on the asset's realized volatility by estimating separately γ^- and γ^+. Given the widely documented negative relationship between changes in volatility and past returns, both coefficients are expected to be negative in order to reflect the well-observed volatility increases in down markets and volatility decreases in up markets. Moreover, an asymmetric volatility effect should be reflected by the coefficient γ^- of negative returns being higher in absolute terms compared to the coefficient γ^+ of positive returns, so that volatility increases after depreciations outweigh on average volatility decreases that follow appreciations of the same magnitude. Finally, the t-statistics across all asymmetry regressions are calculated on the basis of Newey-West (1987) Heteroscedasticity and Autocorrelation Consistent (HAC) standard errors.

Panel B of Table 2.3 presents summary results of estimating equation (2.10) for the index and for the 37 individual stocks. The t-statistics for the index's estimates are reported in parentheses and, for brevity, only the estimates of the γ coefficients are reported. As can be seen from the Table, in the case of the DJIA both gammas are negative and statistically significant, with γ^- equal to -0.868 (t-stat = -6.94) and γ^+ equal to -0.668 (t-stat = -5.44). The above negative values of the signed returns' slopes confirm the predicted negative relationship between past returns and changes in volatility at the daily level and are in line with existing empirical findings. More importantly, though, past negative returns are found to have a greater impact on subsequent index volatility compared to past positive returns, as evidenced by γ^- being higher (in absolute terms) than γ^+, supporting the existence of an asymmetric volatility effect in index returns.

When the regression model in (2.10) is estimated for the 37 individual stocks, the results differ significantly compared to those obtained for the index. On the one hand, the vast majority of the estimated gamma coefficients are negative, confirming the predicted negative relationship between changes in realized volatility and past returns (27 γ^- and 36 γ^+ out of a total of 37 are found to have a negative sign). However, a closer examination of the regression results for the DJIA constituents suggests that volatility asymmetry is far from a uniform characteristic among the sample stocks. The volatility asymmetry present in index returns is not observed for most of the index's components, with only 13 out of the 37 stocks exhibiting lower (more negative) γ^- compared to γ^+. Moreover, the coefficients of interest for negative and for positive returns are -0.2444 (t-stat = -1.09) and -0.5315 (t-stat = -2.22), respectively, for the median stock, i.e. negative returns are found to have on average a smaller impact on subsequent volatility changes compared to positive returns, suggesting an opposite kind of asymmetry in the volatilities of the DJIA components, while similar conclusions can be reached if the coefficients of the lower (L_q) and upper (U_q) quartile stocks are examined. In other words, the specific characteristics of the index's volatility process that refer to a negative correlation between the sum of intraday squared returns and lagged returns, as well as the additional response of volatility to past returns of negative sign, are confirmed for only one third (roughly) of the individual stocks that are included in the index.

Equation (2.11) represents an alternative way of testing for asymmetries in the volatility process of the DJIA and of its constituents. Within this framework, changes in realized volatility are regressed on contemporaneous and past returns, as well as absolute returns. The coefficients $\gamma_{|0|}$ and $\gamma_{|-1|}$ of absolute returns capture the effect of past and of contemporaneous returns, irrespective of their sign. On the other hand, the coefficients γ_0 and γ_{-1} measure the additional impact of negative returns on changes in the realized volatility of the asset.

$$\Delta RV_t = \alpha + \sum_{i=1}^{5} \beta_i \Delta RV_{t-i} + \gamma_0 R_t + \gamma_{-1} R_{t-1} + \gamma_{|0|} |R_t| + \gamma_{|-1|} |R_{t-1}| + \varepsilon_t \tag{2.11}$$

The results from estimating equation (2.11) are reported in Panel C of Table 2.3. The estimated gammas for the index are both negative, indicating a negative correlation between changes in volatility and past/contemporaneous returns. Moreover, the effect of lagged returns, as measured by γ_{-1}, is smaller in magnitude (slope is -0.7463) compared to the effect of contemporaneous returns, measured by γ_0 (-0.7826). It appears, thus, that price movements in the previous trading day have, on average, a lower influence on the magnitude of volatility changes compared to concurrent price movements. Note that both coefficients are statistically significant at any reasonable level, with t-statistics of -11.71 and -10.84 for γ_0 and γ_{-1}, respectively.

The coefficient $\gamma_{|0|}$ of absolute contemporaneous returns for the index is positive (slope is 0.6431) and statistically significant (t-stat = 9.79), while the coefficient $\gamma_{|-1|}$ of absolute lagged returns is positive (slope is 0.0230) but statistically indistinguishable from zero (t-stat = 0.23). Overall, changes in index volatility appear to be negatively correlated with past as well as with contemporaneous returns, with concurrent returns having a greater effect on volatility changes compared to lagged returns. However, given the statistical significance of γ_0 and γ_{-1}, which measure the additional response of volatility to 'bad news', the above negative relationship is more pronounced for negative returns compared to positive ones of similar magnitude. These results provide some additional support for the hypothesis of asymmetric volatility in the time-series of DJIA returns.

When equation (2.11) is estimated for the individual components of the DJIA, the results are relatively in line with those obtained for the index. With respect to absolute returns, $\gamma_{|0|}$ is significantly positive for all 37 individual stocks, although the coefficient $\gamma_{|-1|}$ of absolute lagged returns is consistently negative across all stocks. Moreover, the additional impact of negative returns, as measured by γ_0 and γ_{-1}, is negative and statistically significant for 22 and 23 stocks, respectively.

Summing up, the above results show that the return-generating processes of the index and of its constituents differ substantially, with arguably the most importance difference being the increased level of volatility asymmetry in the case of the index. Section 2.6 investigates whether this difference can be explained by the dynamics of the average realized correlation among the index's components.

2.4.3 Asymmetric GARCH Model

In the class of GARCH models that have been developed to account for the asymmetric responses of volatility to past positive and negative returns, this study uses the extended version of the GJR(1,1) model by Glosten *et al* (1993). The conditional mean and conditional variance of the standard GJR(1,1) specification are given as follows:

$$\begin{aligned} R_t &= \mu + \sqrt{h_t}\, z_t \qquad\qquad z_t \sim i.i.d.N(0,1) \\ h_t &= a_0 + a_1 R_{t-1}^2 + a_2 s_{t-1} R_{t-1}^2 + \beta h_{t-1} \end{aligned} \tag{2.12}$$

where R_t is the asset's return at time t, h_t is the conditional variance at t, the standardized residuals z_t are *independent identically distributed* (i.i.d.) with mean 0 and variance 1, and s_{t-1} is a dummy variable that takes the value of one if the lagged return R_{t-1} is negative and of zero otherwise. The leverage effect is incorporated in the above specification by the use of two coefficients, α_1 and α_2, that separate the effects of past positive and negative returns. More specifically, the impact of lagged positive returns on the conditional variance is given by α_1, while $\alpha_1 + \alpha_2$ provides the impact of lagged negative returns. Since the effect of past returns on volatility is proportional to α_1 and $\alpha_1 + \alpha_2$ for $R_{t-1} \geq 0$

and $R_{t-1} < 0$, respectively, a significantly positive α_2 indicates the existence of a leverage effect, i.e. the conditional variance increases by more following a negative return compared to a positive one. Within this framework, a higher α_2 represents a more pronounced asymmetry in the asset's volatility process.

Following Blair *et al* (2002), the degree of asymmetry in the conditional variance is examined using two measures. The first is termed the 'asymmetry ratio' and is calculated as the ratio of α_2 to the sum $(\alpha_1 + \alpha_2)$. The second measure is the correlation between the stochastic increments in the price and in the variance equation described in Duan (1997). Moreover, assuming symmetrically distributed returns, volatility persistence is estimated as $\alpha_1 + \alpha_2/2 + \beta$. Finally, the GJR specification is estimated using the maximum log-likelihood method (ML) with a Gaussian likelihood.[2]

$$\text{Persistence} = a_1 + \frac{1}{2}a_2 + \beta \tag{2.13}$$

$$\text{Asymmetry Ratio} = \frac{a_1}{a_1 + a_2} \tag{2.14}$$

$$\text{Correlation} = -\frac{a_2}{[\pi(a_1^2 + a_1 a_2 + \frac{5}{8}a_2^2)]^{1/2}} \tag{2.15}$$

The remaining of this Section discusses the results from fitting the standard GJR(1,1) model described in equation (2.12) for DJIA returns and for the returns of its constituents. Table 2.4 reports the point estimates of the GJR specification for the index, as well as for the median, lower quartile and upper quartile stocks. Estimates of volatility persistence and asymmetry are also provided in the last three rows of the Table.

[2] Blair *et al* (2002) suggest that '...while it is unlikely that the returns are generated by a conditional Gaussian distribution, the use of the Gaussian likelihood does ensure consistent estimates of the parameters, which compensates for a loss of efficiency'.

Table 2.4

Estimates of Volatility Parameters for the DJIA and its Constituents using the standard GJR model

$$h_t = a_0 + a_1 R_{t-1}^2 + a_2 s_{t-1} R_{t-1}^2 + \beta h_{t-1}$$

	Index	Stocks			
		Median	L_q	U_q	Greater than Index
α_1	-0.0026 (0.84)	0.0149	0.0097	0.0318	35
α_2	0.1181 (0.00)	0.0550	0.0339	0.0790	3
β	0.9362 (0.00)	0.9416	0.9231	0.9577	19
Persistence	0.9927	0.9976	0.9825	0.9998	25
Correlation	-0.73	-0.56	-0.62	-0.41	34
Ratio	-0.02	0.23	0.14	0.42	35

This Table tabulates the volatility estimates of a standard GJR(1,1) model on the DJIA and its constituent stocks. The sample period is 30 March 1998 to 30 March 2007. Correlation and ratio are measures of volatility asymmetry. P-values for the estimates on the DJIA are in parentheses. The standard errors of the conditional variance's equation are assumed to follow a Student's t distribution.

The α_1 point estimate for the index is -0.0026, compared to 0.0149 for the median stock. Out of the 37 individual stocks in the sample, only one (namely HD) has a α_1 coefficient equal to zero, with the point estimate for the remaining 36 stocks averaging 0.05. With respect to α_2, the point estimate in the case of the DJIA is 0.1181, while the respective value for the median stock is 0.0550. Moreover, α_2 for the upper quartile stock (U_q) is 0.08, and is higher than that of the index for only 3 out of the 37 sample firms. This parameter essentially captures the additional response of volatility to negative innovations and, as has been already discussed, is directly related to asymmetric volatility. Therefore, although α_1 and α_2 need to be considered jointly, the above results seem to indicate that the asymmetric volatility in index returns is significantly more pronounced compared to its constituent stocks.

Overall, the average sensitivity of volatility to past returns, as given by $\alpha_1 + \frac{\alpha_2}{2}$, is slightly higher for the index (0.0578) than for the median stock (0.0531), a result that is mostly attributed to the large additional response of the index's conditional variance to 'bad news', as captured by the high value of α_2. Also, given the constraints of the GJR optimization, it is possible for the estimate of α_2 to be negative as long as the relationship

$\alpha_1 + \alpha_2 \geq 0$ is satisfied. Within the sample of 37 stocks, the estimated α_2 was negative for two companies (namely EK and PFE), but statistically significant for only one of them (PFE).

The coefficient for the lagged variance in index returns is 0.94 and it is equal to the respective value for the median stock. Furthermore, since the point estimates for β in the low and upper quartile stocks are 0.92 and 0.96, respectively, it can be argued that any differences in the volatility processes of the index and of its constituents are mainly attributed to responses to news, as captured by α_1 and α_2. Finally, it should be noted that the estimated GJR parameters satisfy the stationarity constraint $\alpha_0 + \alpha_1 + \frac{1}{2}\alpha_2 + \beta < 1$ for the index and as well as for 35 stocks.

As has been previously mentioned, the asymmetry ratio and the correlation between the stochastic increments in the price and the volatility processes can be used as measures of asymmetric responses of the conditional variance to past returns of different signs. It can be seen from Table 2.4 that both these measures indicate a more pronounced volatility asymmetry in the case of the index compared to its constituent stocks. More specifically, the asymmetry ratio for the index, defined as $\frac{a_1}{a_1 + a_2}$, is equal to -0.02 due to the fact that the α_1 estimate for the DJIA is slightly negative, while the respective value for the median (lower quartile) stock is 0.23 (0.14). Since the ratio includes the total response of volatility to 'bad news' in the denominator, lower values are associated with a more significant leverage effect. Therefore, the fact that 35 out of the 37 companies exhibit higher ratios than that of the index suggests that asymmetric volatility is significantly more pronounced for the index than for its components.

In addition, the correlation measure for the DJIA is -0.73. Given that a more negative correlation indicates a more asymmetric response of the conditional variance to past positive and negative returns, and that only three stocks exhibit higher (absolute) correlation, it is concluded that, within this sample, the 'leverage effect' is more evident in the index compared to individual stocks.

Persistence measures how quickly volatility returns to previous levels after a large positive or negative change or, in other words, how quickly volatility shocks decay away.

Estimated as $\alpha_1 + \frac{1}{2}\alpha_2 + \beta$, higher levels of persistence are associated with long-memory in the conditional variance, i.e. with periods of high (low) volatility lasting longer compared to lower persistence estimates. Based on estimating the GJR(1,1) specification for the DJIA returns, the index exhibits a very high level of volatility persistence (0.9939) while the majority of stocks (25) have even higher point estimates than the index. The median stock has an estimate of 0.9976 and the upper quartile estimate is 0.9998. Overall, only two stocks, namely AXP and TRV, have persistence estimates exactly equal to one (to the sixth decimal point).

2.5 Analysis of the DJIA Average Realized Correlation

This Section discusses the statistical properties of the average realized correlation for the Dow Jones Industrial Average that has been estimated using the methodology described in Section 2.3. Descriptive statistics for the time-series of ρ_{ind} are presented and contrasted with the evolution of the DJIA's realized volatility *RV*. In addition to the univariate distributions of ρ_{ind} and *RV*, the relationship between the two variables is also examined as a preliminary analysis of the hypothesis of correlation dynamics driving the 'asymmetric volatility' phenomenon in DJIA index returns.

2.5.1 Time-Series Properties

Panel A of Table 2.5 presents descriptive statistics for the univariate distributions of ρ_{ind} and *RV*, as well as for their daily changes. As can be seen from the Table, the average level of average realized correlation is 0.17, a value that is relatively close to the median level of 0.16. During the sample period ρ_{ind} ranged from a minimum of 0.01 to a maximum of 0.53, while the standard deviation across the entire period was 0.08. In addition to the mean being slightly higher than the median, positive skewness (0.81) indicates that the distribution of ρ_{ind} is dominated by larger observations in the right side of the distribution, albeit to a relatively small degree. With respect to the fourth moment, the ρ_{ind} distribution is found to be leptokurtic with a kurtosis estimate of 4.01. Finally, the

Jarque-Bera test rejects the hypothesis of normality in the time-series of average realized correlation (t-stat = 355).

Table 2.5

Summary Statistics for the DJIA Average Realized Correlation and Realized Volatility

Panel A: Descriptive Statistics								
	Mean	Median	Min	Max	Standard Deviation	Skewness	Kurtosis	JB Statistic
ρ_{ind}	0.17	0.16	0.01	0.53	0.08	0.81	4.01	355
ρ_{ind} Daily Changes	0.00	0.00	-0.36	0.36	0.07	0.05	4.36	176
RV	0.13	0.12	0.03	0.50	0.06	1.89	8.21	3,906
RV Daily Changes	0.00	0.00	-0.24	0.25	0.04	-0.04	8.37	2,724

Panel B: Autocorrelations					
	lag(1)	lag(2)	lag(3)	lag(4)	lag(5)
ρ_{ind}	0.6106	0.5619	0.5148	0.5136	0.4918
ρ_{ind} Daily Changes	-0.4375	-0.0026	-0.0584	0.0259	0.0017
RV	0.7571	0.7183	0.6746	0.6708	0.6444
RV Daily Changes	-0.4200	0.0099	-0.0821	0.0466	-0.0285

Panel: Cross Correlations							
	-3	-2	-1	0	1	2	3
	DJIA Daily Log-Returns						
ρ_{ind} Daily Changes	0.0254	0.0457	-0.0455	-0.1580	0.0071	0.0126	0.0252
	RV Daily Changes						
ρ_{ind} Daily Changes	-0.0423	-0.0217	-0.2514	0.6086	-0.2494	0.0092	-0.0688

This table tabulates summary statistics for the time series of DJIA average realized correlation (ρ_{ind}) and realized volatility (RV). Panel A provides descriptive statistics for ρ_{ind} and RV as well as for their respective daily changes. Panel B presents the autocorrelations of the series for up to 5 lags. Panel C presents the cross-correlations of index daily log-returns and of daily changes in RV with respect to daily changes in ρ_{ind}, for up to 3 lags and 3 leads.

The average and median daily changes $\Delta\rho_{ind}$ in the index's average realized correlation are both equal to zero, suggesting that there is no trend in the time-series of the average index correlation. The first difference $\Delta\rho_{ind}$ ranges between -0.36 and 0.36, and the standard deviation is 0.07. The distribution of daily correlation changes is almost zero-skewed (skewness = 0.05) but leptokurtic (kurtosis = 4.36). Similarly to the time-series of ρ_{ind}, the Jarque-Bera test rejects normality at any reasonable significance level (t-stat = 176).

The Dow Jones realized volatility RV has a mean of 0.13 per annum across the period from March 1998 to March 2007. RV ranges from a minimum of 0.03 to a maximum of 0.50, and the standard deviation is 0.06. Similarly to the index's realized correlation, the distribution of realized volatility is positively skewed (skewness = 1.89) and highly leptokurtic (kurtosis = 8.21), while the hypothesis of normality is rejected based on the Jarque-Bera t-statistic of 3,906. The first difference ΔRV of realized index volatility has both the mean and median equal to zero, and ranges from -0.24 to 0.25. Contrary to $\Delta\rho_{ind}$, ΔRV exhibits negative skewness (-0.04), but is also significantly leptokurtic (kurtosis = 8.37).

2.5.2 Correlation Persistence

A substantial number of previous studies has found a pronounced long-run dependence in return volatility (see for instance Andersen and Bollerslev (1997b), French, Schwert and Stambaugh (1987) and Blair, Poon and Taylor (2002)), while a similar finding for return correlation has been documented in Andersen *et al* (2001a). Consistent with previous findings, DJIA realized volatility and average realized correlation are found to be indeed characterized by long memory.

Figure 2.2

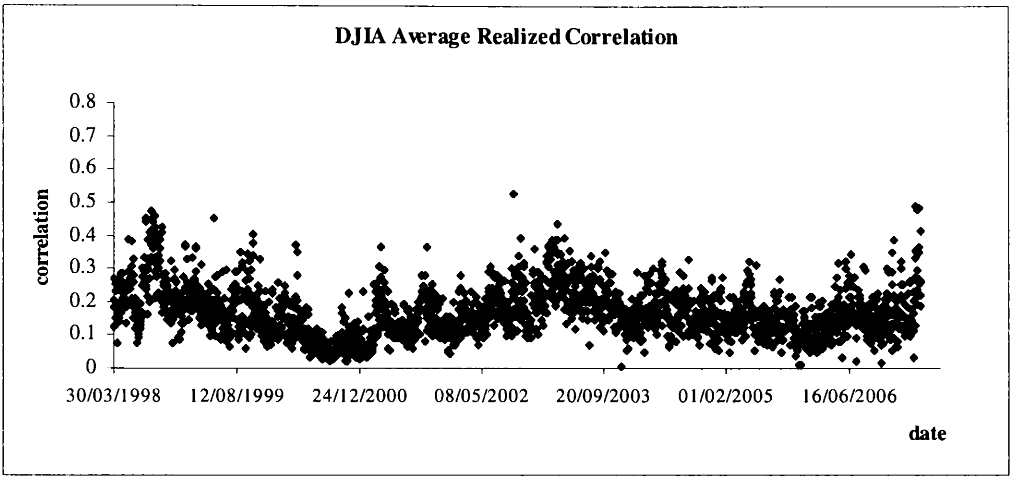

Panel B of Table 2.5 tabulates the autocorrelation coefficients of ρ_{ind} and RV, as well as of their daily changes, for up to five lags. It can be easily seen that autocorrelations in the time series of ρ_{ind} and RV are relatively high at the first five lags, indicating some correlation and volatility clustering. The above clustering is also supported by Figure 2.2 which presents the time evolution of ρ_{ind}, where it can be easily seen that periods of high (low) correlations are more likely to be followed by periods of high (low) correlations. Figure 2.3 plots the time-series of RV and, similarly, shows that periods of high (low) volatilities are systematically followed by periods of high (low) volatilities. Finally, the first differences of the above series, namely $\Delta\rho_{ind}$ and ΔRV, do not exhibit similar long-run dependence, with autocorrelations after the first lag being relatively low in absolute terms. However, autocorrelations at the first lag are significant for $\Delta\rho_{ind}$ and ΔRV (-0.44 and -0.42, respectively).

Figure 2.3

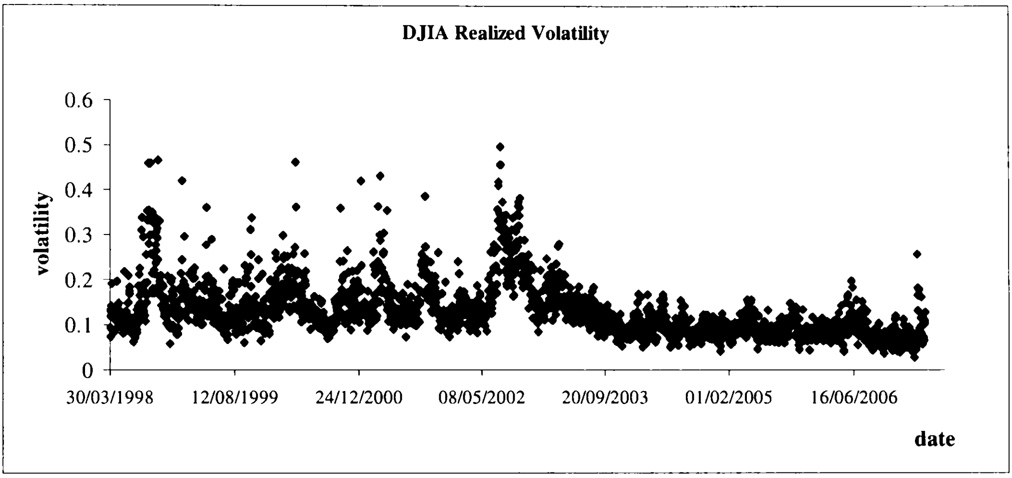

Figure 2.4

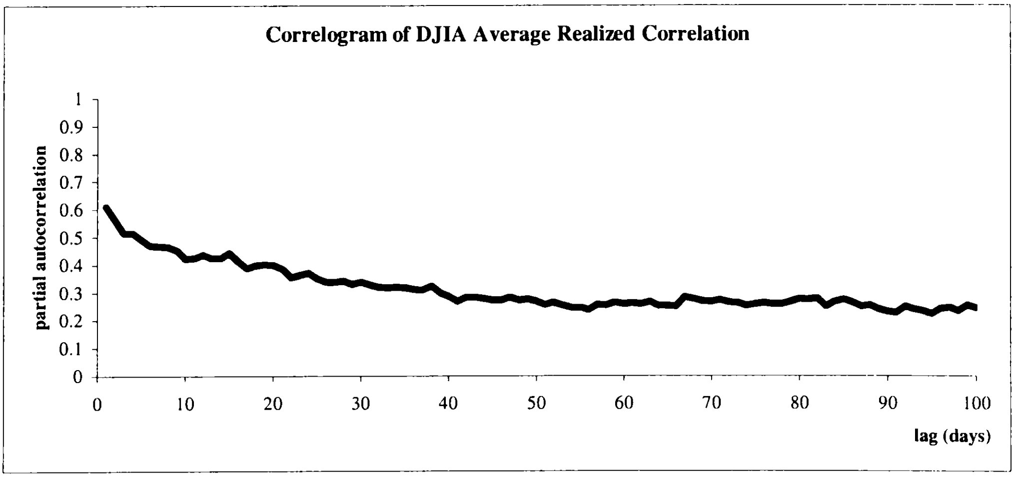

Figure 2.5

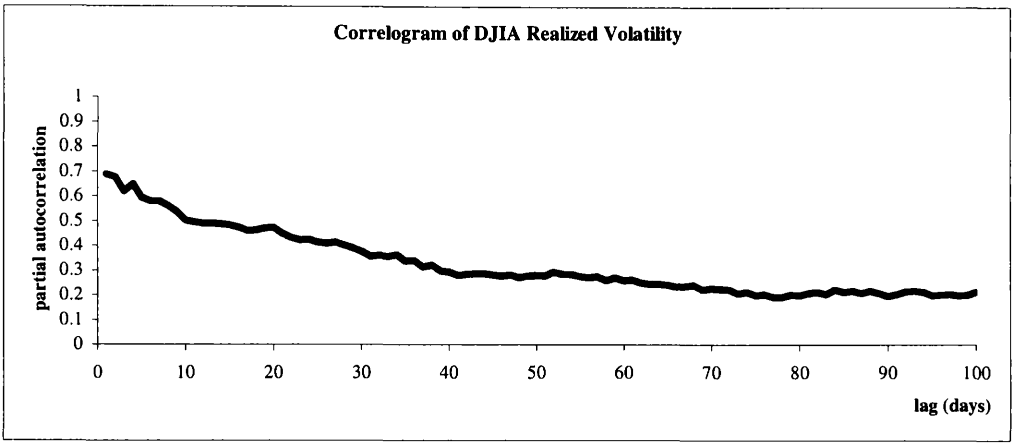

2.5.3 Co-movement of Correlation and Volatility

Andersen *et al* (2001a) document a positive intertemporal relationship between realized correlation and volatility in the Dow Jones index (see also Skintzi and Refenes (2005) for a similar finding using option-implied estimates of DJIA volatility and correlation). This sub-section explores the relationship between ρ_{ind} and RV for the sample period from March 1998 to March 2007.

The scatterplot in Figure 2.6 plots daily changes in the DJIA average realized correlation on daily changes in the index's realized volatility. A positive relationship is evident, indicating that correlations tend to increase (decrease) during periods when volatilities also increase (decrease). This positive relationship between $\Delta\rho_{ind}$ and ΔRV is further explored by regressing daily changes in correlation on contemporaneous and past changes in volatility, controlling for lagged innovations in correlation. The regression model is described in equation (2.16) and the results are presented in Panel A of Table 2.6.

$$\Delta\rho_{ind,t} = \alpha + \sum_{i=1}^{5} \beta_i \Delta\rho_{ind,t-i} + \sum_{j=0}^{5} \gamma_j \Delta RV_{t-j} + \varepsilon_t \qquad (2.16)$$

Figure 2.6

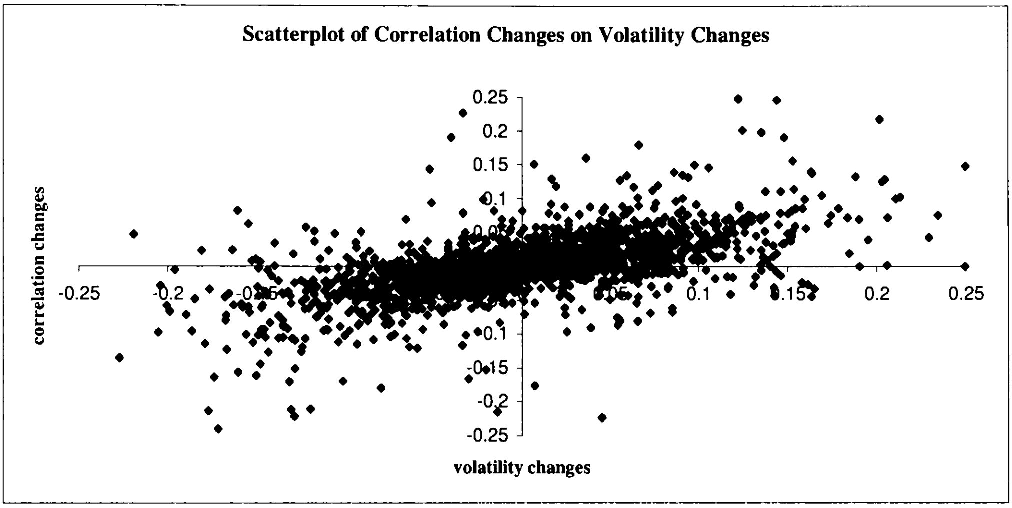

The results from the above specification support the positive relationship between changes in correlation and contemporaneous changes in volatility. The slope coefficient of ΔRV on contemporaneous changes in $\Delta\rho_{ind}$ is 1.1036 and highly significant (t-stat = 49.32). Moreover, past changes in volatility also appear to have an impact on correlation changes, with the coefficients for the first five lags of ΔRV being positive and statistically significant.

The relationship between $\Delta\rho_{ind}$ and ΔRV is further investigated by performing a Granger causality test. Given that a positive relationship between changes in correlation and changes in volatility was documented in the previous analysis, the following test attempts to determine which one is the driving variable behind this positive co-movement. Therefore, it is tested whether lagged changes in RV cause changes in ρ_{ind} and/or lagged changes in ρ_{ind} cause changes in RV. The Granger causality test involves the estimation of the following two regressions:

$$\Delta\rho_{ind,t} = \alpha_1 + \sum_{i=1}^{n} \beta_{1i}\Delta RV_{t-i} + \sum_{i=1}^{n} \gamma_{1i}\Delta\rho_{ind,t-i} + \varepsilon_{1t} \tag{2.17}$$

$$\Delta RV_t = \alpha_2 + \sum_{i=1}^{n} \beta_{2i}\Delta\rho_{ind,t-i} + \sum_{i=1}^{n} \gamma_{2i}\Delta RV_{t-i} + \varepsilon_{2t} \tag{2.18}$$

The first specification tests the null hypothesis H_0: $\beta_{11} = \ldots = \beta_{1n} = 0$ against the alternative that lagged ΔRV Granger cause $\Delta\rho_{ind}$. The second specification tests the null hypothesis H_0: $\beta_{21} = \ldots = \beta_{2n} = 0$ against the alternative that lagged $\Delta\rho_{ind}$ Granger cause ΔRV. Using the Schwarz criterion, the number of lags has been set to n = 5.

Table 2.7 reports the results of the Granger test for the above equations. The F-statistic is 1.42 for equation (2.17), rejecting the null hypothesis of changes in volatility driving changes in correlation. In addition, the F-statistic for equation (2.18) is 121.46, which by far exceeds the critical value (4.74) at a significance level of 5%. Overall, these results provide some initial support for changes in the average correlation in the index's constituents Granger causing contemporaneous changes in the index's realized volatility, and are in line with the 'diversification effect' explanation of asymmetric index volatility.

Table 2.6

Intertemporal Relationship between $\Delta\rho_{ind}$, ΔRV and DJIA Log-Returns

Panel A: Regression of $\Delta\rho_{ind}$ on ΔRV

$$\Delta\rho_{ind,t} = \alpha + \sum_{i=1}^{5} \beta_i \Delta\rho_{ind,t-i} + \sum_{j=0}^{5} \gamma_j \Delta RV_{t-j} + \varepsilon_t$$

α_0	γ_0	γ_1	γ_2	γ_3	γ_4	γ_5	Adj.R^2
0.0001	0.9299 (37.10)	0.5614 (16.24)	0.3040 (8.14)	0.1804 (4.78)	0.0598 (1.64)	0.1935 (6.11)	0.58

Panel B: Regression of $\Delta\rho_{ind}$ on DJIA Log-Returns

$$\Delta\rho_{ind,t} = \alpha + \sum_{i=1}^{5} \beta_i \Delta\rho_{ind,t-i} + \sum_{j=-2}^{0} \gamma_j R_{t+j} + \sum_{j=-2}^{0} \gamma_{|j|} \left|R_{t+j}\right| + \varepsilon_t$$

α_0	γ_0	γ_{-1}	γ_{-2}	$\gamma_{\lvert 0\rvert}$	$\gamma_{\lvert -1\rvert}$	$\gamma_{\lvert -2\rvert}$	Adj.R^2
-0.0019	-0.8681 (-8.28)	-0.8556 (-8.01)	-0.3119 (-2.87)	0.2464 (1.56)	0.1253 (0.81)	-0.0834 (-0.53)	0.34

Panel C: Regression of $\Delta\rho_{ind}$ on DJIA Log-Returns

$$\Delta\rho_{ind,t} = \alpha + \sum_{i=1}^{5} \beta_i \Delta\rho_{ind,t-i} + \gamma^+ D_t^+ R_t + \gamma^- D_t^- R_t$$

α_0	γ^+	γ^-	Adj.R^2
-0.003	-0.4522 (-2.41)	-1.2523 (-6.76)	0.32

Table 2.7

Granger Causality Tests for the Daily Changes in ρ_{ind} and RV

$$\Delta\rho_{ind,t} = \alpha_1 + \sum_{i=1}^{n} \beta_{1i}\Delta RV_{t-i} + \sum_{i=1}^{n} \gamma_{1i}\Delta\rho_{ind,t-i} + \varepsilon_{1t}$$

$$\Delta RV_t = \alpha_2 + \sum_{i=1}^{n} \beta_{2i}\Delta\rho_{ind,t-i} + \sum_{i=1}^{n} \gamma_{2i}\Delta RV_{t-i} + \varepsilon_{2t}$$

Null Hypothesis	F-statistic
$H_0 = \beta_{11} = ... = \beta_{1n} = 0$	1.42
$H_0 = \beta_{21} = ... = \beta_{2n} = 0$	121.46

This Table tabulates the results of the Granger causality tests on daily changes in average realized correlation ($\Delta\rho_{ind}$) and daily changes in realized volatility (ΔRV). The first null hypothesis tests whether changes in RV Granger cause changes in ρ_{ind}. The second hypothesis tests whether changes in ρ_{ind} Granger cause changes in RV. The critical value is 4.735 (at the 5% significance level, for 10 degrees of freedom in the numerator and 5 degrees of freedom in the denominator).

2.5.4 Correlation Asymmetry

Skintzi and Refenes (2005) find that DJIA implied correlation responds asymmetrically to past positive and negative index returns, while Andresen *et al* (2001a) reach the same conclusion using measures of realized correlation. Following this line of thought, the regression model in equation (2.19) is estimated, in an attempt to examine the effect of past index returns on the index's average realized correlation in the full sample:

$$\Delta\rho_{ind,t} = \alpha + \sum_{i=1}^{5} \beta_i \Delta\rho_{ind,t-i} + \sum_{j=-2}^{0} \gamma_j R_{t+j} + \sum_{j=-2}^{0} \gamma_{1j1} \left|R_{t+j}\right| + \varepsilon_t \tag{2.19}$$

where R_t is the daily log-return of the DJIA at time t. The first five lags of the $\Delta\rho_{ind}$ series have been included in order to account for the autocorrelation structure of ρ_{ind} described in 2.5.2. The main variables of interest are the contemporaneous and lagged index returns, as well as their absolute values. Since previous empirical findings suggest that correlations between securities increase during and after down markets and decrease when the market rises, one would expect the coefficients γ_0, γ_{-1} and γ_{-2} of lagged index returns to be negative and statistically significant. Furthermore, significantly positive coefficients of absolute index returns ($\gamma_{|j|}$, for j = [0, -1, -2]) would indicate that the above negative relationship is also dependent on the magnitude of past index returns.

Panel B of Table 2.6 reports the regression results for equation (2.19). The coefficient of R_t is negative (γ_0 = -0.8681) and statistically significant (t-stat = -8.28), confirming the predicted negative relationship between changes in correlation and contemporaneous changes in the index level. The fact that coefficients of index returns at the first two lags are also negative (γ_{-1} = -0.8556 and γ_{-2} = -0.3119) and statistically significant (the t-stats are -8.01 and -2.87, respectively) provides further evidence for the hypothesis that correlations among an index's constituents tend to increase when the index drops and to decrease when the index rises. This finding is related to Rubinstein's 'diversification effect' which states that at periods of declining markets, securities tend to move more closely together, thereby reducing the opportunities for diversification and increasing the risk of investing in the index. Moreover, the absolute magnitude of index returns appears to have a small impact on contemporaneous changes in correlation since the $\gamma_{|0|}$ coefficient of $|R_t|$ is found to be positive ($\gamma_{|0|}$ = 0.2464) albeit marginally insignificant (t-stat = 1.56). In other words, $\Delta\rho_{ind}$ is not only affected by the sign of contemporaneous index returns, but it might also be dependent on the return's magnitude, with a larger fall (rise) in the index being associated with a higher increase (decrease) in the correlations among constituent stocks.

After establishing that movements in the DJIA negatively affect changes in contemporaneous and in future correlations, the extent to which correlation responds symmetrically to index returns of different signs is examined. In order to distinguish between changes in correlation following positive and negative returns, the following regression is estimated (see also Skintzi and Refenes (2005)):

$$\Delta\rho_{ind,t} = \alpha + \sum_{i=1}^{5} \beta_i \Delta\rho_{ind,t-i} + \gamma^{+} D_t^{+} R_t + \gamma^{-} D_t^{-} R_t + \varepsilon_t \qquad (2.20)$$

where D_t^{+} is a dummy variable that is equal to one if the contemporaneous index return is positive, and equal to zero otherwise. Accordingly, the dummy variable D_t^{-} takes the value of one if R_t is negative, and of zero otherwise. Within this framework, the coefficient γ^+ refers to the impact of positive index returns on contemporaneous changes in correlation, while γ^- refers to the impact of negative index returns on $\Delta\rho_{ind}$.

As can be seen from Panel C of Table 2.6, similarly to the previous findings in (2.19), the γ^+ and γ^- coefficients in equation (2.20) are both negative and statistically significant (t-stats are -2.41 and -6.76, respectively), supporting the existence of a negative relationship between $\Delta\rho_{ind,t}$ and R_t. However, the coefficient γ^- for negative index returns is significantly larger (in absolute magnitude) compared to γ^+ for positive returns, indicating a pronounced asymmetry in the response of correlation to market movements of different signs. More specifically, given that the coefficient γ^- for negative R_t is almost three times the magnitude of γ^+ for positive index returns (the coefficients are -0.4522 and -1.2523, respectively), it appears that correlation increases during negative market movements are on average higher than correlation decreases during positive market movements by a factor of three. This asymmetric relationship is consistent with findings by Andersen *et al* (2001a) and by Skintzi and Refenes (2005).

2.6 The Impact of Correlation Dynamics on DJIA Volatility Asymmetry

2.6.1 Extended Asymmetric Regressions

The 'diversification effect' suggests that changes in the average index correlation are the main source of asymmetry in the index's volatility. Within this framework, volatility increases (decreases) are caused by past increases (decreases) in the correlations among the index's constituents which, on average, tend to coincide with negative (positive) index returns. The above hypothesis is examined by estimating an extended version of equation (2.10), where lagged changes in the average realized index correlation $\Delta\rho_{ind,t-1}$ are introduced as an additional explanatory variable for changes in the index's volatility ΔRV_t. According to the 'diversification effect' hypothesis, lagged correlation changes are expected to be significantly correlated with volatility changes, thereby decreasing the explanatory power of lagged returns. In addition, given that correlation increases in down markets on average outweigh correlation decreases in up markets (see Section 2.5.4), the effect of past index returns is expected to be similar irrespective of the returns' signs, i.e. $\gamma^- = \gamma^+$. The extended model is given in equation (2.21) and the results are presented in Panel A of Table 2.8.

$$\Delta RV_t = \alpha + \sum_{i=1}^{5} \beta_i \Delta RV_{t-i} + \gamma^- s^- R_{t-1} + \gamma^+ s^+ R_{t-1} + \delta \Delta \rho_{ind,t-1} + \varepsilon_t \tag{2.21}$$

The estimated coefficients of (2.21) do not seem to support the 'diversification' explanation of asymmetric index volatility. More specifically, the coefficients of lagged returns remain relatively unchanged compared to those of the standard model in (2.10), with γ^- = -0.8683 and γ^+ = -0.6762. However, their statistical significance after controlling for correlation dynamics is reduced, as evidenced by t-statistics of -4.29 and of -4.54 for γ^- and γ^+, respectively (compared to t-statistics of -6.94 and -5.44 in equation (2.10)). Furthermore, the difference in the coefficients' magnitudes indicates that the asymmetry in index volatility is as pronounced as in the standard specification, with volatility increases after negative returns being on average higher than volatility decreases after positive returns of similar absolute levels. The failure of $\Delta\rho_{ind,t-1}$ to explain volatility changes is further demonstrated by its coefficient δ being statistically insignificant (t-stat = -0.83). Finally, somewhat surprisingly, the estimated coefficient δ of lagged correlation changes is found to be negative and, thus, in contrast with the theoretical prediction, as well as with previous empirical findings in Section 2.5.3, of a positive co-movement of correlation changes with volatility changes.

An extended version of the asymmetric equation (2.11) is also tested, where correlation changes have been added as an additional independent variable. The new specification is given in (2.22) and the results are tabulated in Panel B of Table 2.8. Similarly to the extended specification in (2.21), the introduction of $\Delta\rho_{ind,t-1}$ is expected to reduce the impact of lagged returns on volatility changes, reflected by less significant returns' coefficients. Moreover, the additional impact of past negative returns, as captured by γ_{-1}, is expected to decrease, or even become insignificant after controlling for correlation dynamics.

$$\Delta RV_t = \alpha + \sum_{i=1}^{5} \beta_i \Delta RV_{t-i} + \gamma_0 R_t + \gamma_{-1} R_{t-1} + \gamma_{|0|} |R_t| + \gamma_{|-1|} |R_{t-1}| + \delta \Delta \rho_{ind,t-1} + \varepsilon_t \qquad (2.22)$$

Similarly to the extended specification in (2.21), the results from estimating (2.22) indicate that correlation dynamics fail to account for the asymmetric volatility of the DJIA. Although the statistical significance of the returns' gammas reduces, their magnitude remains relatively unchanged after including $\Delta\rho_{ind,t-1}$ in the specification. Furthermore, negative index returns still have an additional impact on volatility changes compared to positive ones, while the coefficient of lagged correlation changes is statistically indistinguishable from zero (t-stat = -0.73).

Table 2.8

Asymmetric Volatility Regressions Controlling for Correlation Dynamics

Panel A: $\Delta RV_t = \alpha + \sum_{i=1}^{5} \beta_i \Delta RV_{t-i} + \gamma^- s^- R_{t-1} + \gamma^+ s^+ R_{t-1} + \delta \Delta \rho_{ind,t-1} + \varepsilon_t$

γ^-	γ^+	δ	Adj.R^2
-0.8683	-0.6762	-0.0145	0.30
(-4.29)	(-4.54)	(-0.83)	

Panel B: $\Delta RV_t = \alpha + \sum_{i=1}^{5} \beta_i \Delta RV_{t-i} + \gamma_0 R_t + \gamma_{-1} R_{t-1} + \gamma_{|0|} |R_t| + \gamma_{|-1|} |R_{t-1}| + \delta \Delta \rho_{ind,t-1} + \varepsilon_t$

γ_0	γ_{-1}	$\gamma_{\vert 0\vert}$	$\gamma_{\vert -1\vert}$	δ	Adj.R^2
-0.7809	-0.7498	0.6452	0.0193	-0.0123	0.35
(-7.33)	(-7.61)	(3.89)	(0.13)	(-0.73)	

This Table tabulates the results of regressions examining the asymmetric responses of realized volatility to past returns of opposite signs controlling for the dynamics of the index's average realized correlation. The regressions are estimated on the returns of the DJIA for the period 30 March 1998 to 30 March 2007

2.6.2 Extended GJR-C Model

In order to directly examine whether correlation dynamics have an impact on the evolution of h_t, an additional term $\Delta\rho_{ind,t-1}$ is introduced into the equation for the conditional variance of the GJR specification in (2.12). The variable $\Delta\rho_{ind,t-1}$ is a measure of lagged changes in the average correlation among the index's constituents using high-frequency data on the index and on its constituents, estimated using the Andersen *et al*

(2001b) methodology described in the previous subsections. The extended specification is termed GJR(1,1)-C and is given as follows:

$$
\begin{aligned}
&R_t = \mu + \sqrt{h_t}\, z_t \qquad\qquad z_t \sim i.i.d.N(0,1) \\
&h_t = a_0 + a_1 R_{t-1}^2 + a_2 s_{t-1} R_{t-1}^2 + \beta h_{t-1} + \gamma \Delta\rho_{ind,t-1}
\end{aligned}
\tag{2.23}
$$

The hypothesis of correlation changes driving the leverage effect in index volatility predicts that, after controlling for correlation dynamics in the extended GJR-C model, α_2 should decrease compared to the estimate in the standard GJR model, or even become statistically indistinguishable from zero. An insignificant α_2 would indicate that the conditional variance responds symmetrically to past returns, irrespective of their sign, and that only the absolute magnitude of R_{t-1} has any incremental effect on index volatility in addition to that of changes in the average correlation among the index's constituents.

Panel A of Table 2.9 reports the results of the extended GJR-C model for the returns of the Dow Jones Industrial Average. The extended specification includes the lagged daily changes $\Delta\rho_{ind,t-1}$ in the index's average realized correlation as an exogenous variable in the equation of the index's conditional variance, in an attempt to control for correlation dynamics when examining the leverage effect in DJIA index returns. In addition to the estimated parameters, the Table tabulates standard errors of fit as well as the corresponding z-statistics and p-values. Variance persistence and the two measures of asymmetry previously described are also presented.

As can be seen from Panel A, when equation (2.23) is estimated, α_1 slightly decreases from -0.0026 to -0.0040 and remains statistically insignificant (p-value is 0.67). More importantly, incorporating correlation changes does not appear to have any impact on the asymmetric responses of the conditional variance to lagged returns of opposite signs. This is reflected in the estimate α_2 of lagged negative returns which, instead of decreasing as the 'diversification' hypothesis predicts, experiences a small increase from 0.1181 to 0.1201. Furthermore, the coefficient γ for the lagged changes in correlation is statistically indistinguishable from zero at the 5% confidence level, and only marginally insignificant at the 10% level (p-value is 0.1071).

Table 2.9

Estimates of Volatility Parameters for the DJIA using the extended GJR-C model

Panel A: $h_t = a_0 + a_1 R_{t-1}^2 + a_2 s_{t-1} R_{t-1}^2 + \beta h_{t-1} + \gamma \Delta\rho_{ind,t-1}$

	Coefficient	St. Error	z-statistic	p-value
α_0	0.0000	0.0000	4.13	0.0000
α_1	-0.0040	0.0093	-0.43	0.6681
α_2	0.1201	0.0145	8.27	0.0000
β	0.9367	0.0094	99.82	0.0000
γ	0.0000	0.0000	-1.61	0.1071
Persistence	0.9927			
Correlation	-0.73			
Ratio	-0.03			

Panel B: $h_t = a_0 + a_1 R_{t-1}^2 + a_2 s_{t-1} R_{t-1}^2 + \beta h_{t-1} + \gamma_1 \Delta\rho_{ind,t-1} + \gamma_2 s_{t-1} \Delta\rho_{ind,t-1}$

	Coefficient	St. Error	z-statistic	p-value
α_0	0.0000	0.0000	0.46	0.6462
α_1	0.0067	0.0106	0.63	0.5275
α_2	0.0864	0.0153	5.64	0.0000
β	0.9375	0.0094	99.85	0.0000
γ_1	0.0000	0.0000	1.59	0.1117
γ_2	0.0000	0.0000	-4.13	0.0000
Persistence	0.9874			
Correlation	-0.67			
Ratio	0.07			

This Table tabulates the volatility estimates of the modified GJR-C model on the returns of the DJIA. The sample period is 30 March 1998 to 30 March 2007. Panel A presents the simple GJR-C model of unconditional correlation changes, while Panel B presents the extended GJR-C model of conditional correlation changes. 'Correlation' and 'ratio' are measures of volatility asymmetry. P-values for the estimates on the DJIA are in parentheses. The standard errors of the conditional variance's equation are assumed to follow a Student's t distribution.

This inability of unconditional correlation dynamics to account for the index's asymmetric volatility is also highlighted by the two asymmetry measures, namely correlation and the asymmetry ratio, the estimates of which are similar to those for the standard GJR specification in (2.12). Finally, the Likelihood Ratio Test (LRT) rejects the hypothesis that adding unconditional correlation changes as an exogenous regressor in the GJR's variance equation results in a significantly better fit across the sample data.

The LTR compares the likelihood score of the extended GJR-C in (2.23) to that of the standard GJR in (2.12), with the latter specification considered as a nested version of the former, by computing the following Likelihood Ratio (*LR*):

$$LR = 2(L1 - L2) \tag{2.24}$$

where *L1* and *L2* refer to the Likelihood scores of the GJR-C and of the GJR specifications, respectively. Given that the GJR model is a nested version of the GJR-C, *L1* will be higher than *L2* merely due to the use of an additional explanatory variable. However, whether using an additional variable is justified by a significantly better fit can be determined by comparing *LR* with a critical value under the assumption that the ratio approximately follows a chi-square distribution. Across the DJIA sample in this study, the above *LR* is equal to 1.66. Since the critical value at the 5% confidence level and with one degree of freedom is 3.84, the relatively low *LR* indicates that incorporating $\Delta\rho_{ind,t-1}$ in the specification does not result in a significantly better fit compared to the standard GJR.

The above results seem to cast some doubt on the hypothesis of correlation changes driving asymmetric volatility in index returns. In addition to the slight increase of the coefficient for the additional response of index variance to negative returns, correlation innovations are not found to be significantly correlated with the index's conditional variance. This finding stands in contrast to the previously reported positive comovement between realized index volatility *RV* and average realized correlation ρ_{ind}, and does not support the 'diversification effect' explanation of asymmetric index volatility. This initial rejection of the 'diversification' hypothesis is particularly puzzling given that correlation changes were found to be negatively correlated with index returns, and that this relationship was stronger for negative returns compared to positive ones.

In addition to the general specification in (2.23), an alternative version of the GJR-C model is examined, in which the relationship between the conditional variance and changes in the index's average correlation is conditional upon the sign of past index returns, as given by the s_{t-1} dummy. The specification in (2.25) effectively explores whether the impact of correlation changes on the index's variance depends on the direction of index movements. Similarly to distinguishing between the impact of past positive and negative returns through the pair of coefficients α_1 and α_2, the coefficients γ_1 and γ_2 measure the effect of correlation changes on the index's conditional variance in the case of up and of down markets, respectively.

$$h_t = a_0 + a_1 R_{t-1}^2 + a_2 s_{t-1} R_{t-1}^2 + \beta h_{t-1} + \gamma_1 \Delta\rho_{ind,t-1} + \gamma_2 s_{t-1} \Delta\rho_{ind,t-1} \quad (2.25)$$

Panel B of Table 2.9 reports the results for the above version of the GJR-C. When the effect of correlation changes is combined with the use of the dummy s_{t-1} for the sign of past index returns, the intercept α_0 remains close to zero but becomes statistically insignificant (p-value is 0.67), while the coefficient β for the first lag of the conditional variance remains at approximately the same level (0.9375). More importantly, though, the lack of explanatory power of $\Delta\rho_{ind,t-1}$ over the DJIA's conditional variance now appears to be limited only to up markets. The coefficient γ_1 of lagged correlation changes conditional on a positive lagged index return is very close to zero and statistically insignificant (p-value is 0.11), similarly to the estimate of γ in (2.23). On the other hand, correlation changes that coincide with a negative index return have a statistically significant impact, measured by γ_2 which, although admittedly small in magnitude, has a p-value of zero.

Furthermore, separating the effect of correlation changes in up and in down markets results in a less pronounced 'asymmetric volatility' phenomenon in the DJIA returns. First, the coefficient α_1 of lagged positive returns experiences a small increase while remaining statistically indistinguishable from zero. More noticeably, the coefficient α_2 of the additional impact of lagged negative returns decreases significantly to 0.0864 compared to the estimate of the standard GJR model (where α_2 = 0.1181) and to that of the extended GJR-C model in (2.23) (where α_2 = 0.1201). This reduction in the asymmetric response of index volatility to past returns of different signs can be further demonstrated by the changes in the 'asymmetry ratio' and the 'correlation' measures. The 'correlation' measure, in particular, increases from -0.73 to -0.67 while the 'asymmetry ratio' increases from -0.03 to 0.07, with higher (or less negative) values for both measures indicating a weaker asymmetry in the index's conditional variance. Finally, the respective Likelihood Ratio LR is 13.73. Given that the critical value at the 5% level and with two degrees of freedom is 5.99, the relatively high LR does not reject the hypothesis that adding conditional correlation changes in the standard specification results in a significantly better fit across the sample data.

Overall, accounting for the dynamics of the DJIA's average realized correlation results in a reduction of the observed asymmetry in the index's volatility process. Not surprisingly, correlation changes are found to be significantly correlated with the index's conditional variance when the index falls, while their comovement is much weaker when the index rises. This asymmetric relationship between correlation changes and volatility is in line with previous empirical studies which suggest that the so-called 'leverage effect' might be more appropriately considered as a 'down-market' effect. Although index volatility tends to increase after negative index returns and, to a lesser extent, decrease after positive returns, the return's magnitude is found to have a somewhat smaller impact on this relationship than that previously documented. Moreover, the inability of unconditional correlation changes to explain the index's volatility process contrasted to the significant coefficient of the conditional variable indicates that correlation dynamics alone cannot substitute the effect of a dummy variable for the sign of market movements within a GARCH framework. However, given that the coefficient α_2 for lagged negative returns decreases after introducing the conditional $\Delta\rho_{ind,t-1}$, correlation changes seem to absorb some of the explanatory power of the magnitude of past returns which, nevertheless, remains significant.

2.7 Conclusion

This Chapter has examined the widely reported asymmetric volatility of index returns from the perspective of the 'diversification effect', using data on the Dow Jones Industrial Average across a ten-year period. When asymmetric GARCH models are fitted on index returns, it is commonly found that the conditional variance experiences an increase after negative returns that is higher compared to decreases following positive returns of similar magnitude. Moreover, this finding is significantly more pronounced for indices than for individual stocks, highlighting a fundamental difference between the volatility processes of these two asset classes. In contrast to previous explanations referring to a 'leverage effect' or to 'volatility feedback', this study explores whether changes in the correlations among the index constituents can account for the observed volatility asymmetry.

A set of initial results is presented that is consistent with the 'diversification effect' hypothesis. More specifically, the average realized correlation of the DJIA is found to be negatively correlated with index returns, confirming that diversification opportunities reduce in down markets where the majority of stocks move more closely together. This relationship is found to be asymmetric, with correlation increases coinciding with negative index returns being higher than correlation decreases during positive returns. Furthermore, the average correlation is positively related to the index's realized volatility, indicating that periods of high (low) volatility are also characterized by high (low) correlations among the constituent stocks.

Within a GARCH framework, evidence is provided that changes in the average correlation among the constituents are directly linked to the asymmetric conditional variance of the parent index. Estimating an extension of the GJR specification, where conditional innovations in the average index correlation are included as an exogenous regressor, results in a lower coefficient of past negative returns compared to that of the standard model. Higher estimates of two asymmetry measures, namely the 'correlation' and the 'asymmetry ratio', also suggest that volatility asymmetry in DJIA index returns is less pronounced after controlling for correlation dynamics.

In addition, correlation changes are significantly related to the conditional variance in down markets and only weakly correlated in up markets. This asymmetric comovement supports previous empirical findings that refer to the 'leverage' effect as a 'down-market' effect, since the direction of past market movements remains highly significant in explaining the index's conditional variance. Therefore, given the inability of unconditional correlation changes to explain the asymmetric volatility in DJIA returns, it appears that conditional correlation changes in fact absorb some of the explanatory power of the magnitude of lagged index returns. However, it has to be noted that accounting for correlation dynamics does not fully explain the asymmetry of the conditional variance, and that the sign as well as the magnitude of past returns are still highly significant in predicting future index volatility.

References

Andersen, Torben, and Tim Bollerslev, 1997a, Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns, *Journal of Finance* 52, 975-1005

Andersen, Torben, and Tim Bollerslev, 1997b, Intraday Periodicity and Volatility Persistence in Financial Markets, Journal *of Empirical Finance* 4, 115-158

Andersen, Torben, Tim Bollerslev, Francis Diebold, and Heiko Ebens, 2001a, The distribution of Realized Stock Return Volatility, *Journal of Financial Economics* 61, 43-76

Andersen, Torben, Tim Bollerslev, Francis Diebold, and Paul Labys, 2001b, Exchange Rate Returns Standardized by Realized Volatility are (nearly) Gaussian, *Journal of the American Statistical Association* 96, 42-55

Bekaert, Geert, and Guojun Wu, 2000, Asymmetric Volatility and Risk in Equity Markets, *Review of Financial Studies* 13, 1-42

Black, Fisher, 1976, Studies of Stock Price Volatility Changes, *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section*, 177-181

Blair, Bevan, Ser-Huang Poon, and Stephen Taylor, 2002, Asymmetric and Crash Effects in Stock Volatility for the S&P 100 Index and its Constituents, *Applied Financial Economics* 12, 319-329

Bollerslev, Tim, Julia Litvinova, and George Tauchen, 2006, Leverage and Volatility Feedback Effects in High-Frequency Data, *Journal of Financial Econometrics* 4, 353-384

Campbell, John, and Ludger Hentschel, 1992, No News is Good News, *Journal of Financial Economics* 31, 281-318

Cappiello, Lorenzo, Robert Engle, and Kevin Sheppard, 2006, Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns, *Journal of Financial Econometrics* 4, 537-572

Cheung, Yin-Wong, and Lilian Ng, 1992, Stock Price Dynamics and Firm Size: An Empirical Investigation, *Journal of Finance* 48, 1985-1997

Christie, Andrew, 1982, The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects, *Journal of Financial Economics* 3, 407-432

Driessen, Joost, Pascal Maenhout, and Grigory Vilkov, 2008, The Price of Correlation Risk: Evidence from Equity Options, forthcoming in the *Journal of Finance*

Duan, Jin-Chuan, 1997, Augmented GARCH(p-q) Process and its Diffusion Limit, *Journal of Econometrics* 79, 97-217

Engle, Robert, and Victor Ng, 1993, Measuring and Testing the Impact of News on Volatility, *Journal of Finance* 48, 1749-1778

Figlewski, Stephen, and Xiaozu Wang, 2000, Is the 'Leverage Effect' a Leverage Effect?, *NYU working paper*

French, Kenneth, G. William Schwert, and Robert Stambaugh, 1987, Expected Stock Returns and Volatility, *Journal of Financial Economics* 19, 3-30

Glosten, Lawrence, Ravi Jagannathan, and David 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

Han, Bing, 2007, Stochastic Volatilities and Correlations of Bond Yields, *Journal of Finance* 62, 1491-1524

Kim, Dogncheol, and Stanley Kon, 1994, Alternative Models for the Conditional Heteroscedasticity of Stock Returns, *Journal of Business* 67, 563-598

Kroner, Kenneth, and Victor Ng, 1998, Modeling Asymmetric Comovements of Asset Returns, *Review of Financial Studies* 11, 817-844

Lamoureux, Christopher, and William Lastrapes, 1990, Persistence in Variance, Structural Change, and the GARCH Model, *Journal of Business and Economic Statistics* 8, 225-234

Nelson, Daniel, 1991, Conditional Heteroscedasticity in Asset Returns: A New Approach, *Econometrica* 59, 347-370

Newey, Whitney, and Kenneth West, 1987, A Simple, Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703-708

Pollet, Joshua, and Mungo Wilson, 2008, Average Correlation and Stock Market Returns, forthcoming in the *Journal of Financial Economics*

Rubinstein, Mark, 2000, Comments on the 1987 Stock Market Crash: 11 Years Later, *Risks in Investment Accumulation Products of Financial Institutions*, published by the Society of Actuaries

Skintzi, Vasiliki, and Apostolos Refenes, 2005, Implied Correlation Index: A New Measure of Diversification, *Journal of Futures Markets* 25, 171-197

Smith, Daniel, 2007, Risk and Return in Stochastic Volatility Models: Volatility Feedback Matters!, *Simon Fraser University working paper*

Stivers, Christopher, Patrick Dennis, and Stewart Mayhew, 2006, Stock Returns, Implied Volatility Innovations and the Implied Volatility Phenomenon, *Journal of Financial and Quantitative Analysis* 41, 381-406

Tauchen, George, Harold Zhang, and Ming Liu, 1996, Volume, Volatility and Leverage: A Dynamic Analysis, *Journal of Econometrics* 74, 177-208

Wu, Guojun, 2001, The Determinants of Asymmetric Volatility, *Review of Financial Studies* 14, 837-859

Chapter 3
The Asymmetric Impact of Firm-Specific and of Index Returns on the Volatility Processes of Individual Stocks

3.1 Introduction

3.1.1 Literature Review

ARCH models have been increasingly popular in describing the returns generating process of stocks and of indices, with a large number of studies reporting that volatility processes tend to be highly persistent at a daily frequency. However, the high volume of related papers has highlighted some significant differences between the volatility processes of individual stocks and of indices. Arguably one of the most important differences refers to the extent to which volatility responds negatively and asymmetrically to past returns, with this well-documented empirical irregularity usually termed the 'volatility asymmetry' phenomenon.

The stylized fact that volatility is negatively correlated with stock returns, and that this correlation is conditional on the return's sign, has been examined by Black (1976) and by Christie (1982). These were among the first papers to show that volatility increases after negative returns tend to outweigh volatility decreases following positive returns of similar magnitude, and both related this asymmetry to changes in leverage. The intuition behind the terms 'leverage effect' and 'volatility asymmetry' being used almost interchangeably in the early literature is that a negative return will result in an increase of a firm's debt-to-equity ratio (i.e. its leverage), and this will be reflected in an increase of the underlying equity risk and a higher equity volatility.

A second stream of the literature has proposed 'volatility feedback' as an alternative explanation for asymmetric volatility. This line of thought reverses the causality by suggesting that an increase in volatility at time t will result in an increased risk-premium at $t+1$ to compensate investors for bearing more risk. Increased risk-premia, though, at $t+1$ are equivalent to a decreased stock price at t, thereby introducing a negative contemporaneous relationship between stock returns and volatility. The volatility feedback explanation has been examined by Campbell and Hentschel (1992), French, Schwert and Stambaugh (1987) and Smith (2007), while comparative analyses of the two hypotheses have been conducted by Bekaert and Wu (2000) and Bollerslev, Litvinova and Tauchen (2006).

Kim and Kon (1994) focus on thirty large-capitalization stocks as well as on three indices, and find that individual stock volatilities are significantly less asymmetric than index volatilities. Similar conclusions have also been reached by Tauchen, Zhang and Liu (1996) and by Blair, Poon and Taylor (2002). This discrepancy between stocks and indices, however, cannot be accommodated by either the 'leverage effect' or the 'volatility feedback' explanations, since the fundamental causes for the asymmetry do not depend on the asset class examined.

The possibility of market-level influences driving the 'asymmetric volatility' phenomenon has been discussed by Figlewski and Wang (2000) who examine a large sample of S&P 100 stocks as well as the index itself. After accounting for changes in the debt-to-equity ratio, they find that leverage changes cannot adequately explain changes in the realized and in the implied volatility of the sample stocks, especially in up markets. The authors then conclude that the 'volatility asymmetry' phenomenon could be more appropriately described as a 'down-market' effect that is not necessarily related to changes in leverage or to volatility feedback.

Finally, the present study is directly related to the findings by Stivers, Dennis and Mayhew (2006). Assuming that implied volatility acts as an observable proxy for expected return volatility, they define the 'asymmetric volatility' phenomenon as the relationship between returns and innovations in implied volatility. Focusing on the US equity market, Stivers, Dennis and Mayhew (2006) report that, although index returns are strongly negatively correlated with changes in index implied (systematic) volatility, individual stock returns exhibit only a weak correlation with changes in firm-specific implied (idiosyncratic) volatility. This could be interpreted as additional evidence of a more pronounced asymmetry effect in the case of indices compared to individual stocks. However, stock returns are found to exhibit a significant negative correlation with innovations in systematic volatility, i.e. in index implied volatility, a result that is consistent with the hypothesis of 'asymmetric volatility' being driven by market factors rather by changes in leverage or by volatility feedback.

3.1.2 Scope of Study

This Chapter examines the 'volatility asymmetry' phenomenon in the returns of individual stocks. Previous empirical papers have reported that the volatility processes of stocks differ significantly from those of indices, particularly in being less asymmetric with respect to past returns of opposite signs. Given that the two most common explanations for asymmetric volatility, namely the 'leverage effect' and 'volatility feedback', are not able to reconcile the above difference, this analysis focuses on the hypothesis that asymmetric volatility is in essence a 'down-market' effect, proposed, among others, by Figlewski and Wang (2000). Within this framework, the possibility is explored that the conditional variance of individual stocks is asymmetrically correlated with past market returns instead of (or in addition to) firm-specific returns.

The sample consists of the returns of the thirty Dow Jones's constituents across the ten-year period from January 1998 to December 2007. When a standard asymmetric GJR model (developed by Glosten, Jagannathan and Runkle (1993)) is estimated, individual stock volatilities are shown to be significantly less asymmetric than the index's volatility, confirming previous empirical findings. In order to directly examine the impact of market-shocks on the stocks' conditional variances, a modification of the GJR is tested, termed GJR-I, where the lagged signed returns of the index have replaced the lagged signed firm-specific returns. First, it is found that individual stock volatilities are significantly correlated with past index returns. Moreover, volatilities respond asymmetrically to returns of opposite signs, with volatility increases after a negative market return being on average higher than volatility decreases after a positive market return of similar magnitude. This asymmetry is generally more pronounced compared to the one observed in the standard GJR model, based on the standard 'asymmetry ratio' measure.

In addition to the standard GJR-I, an extension is estimated where the conditional variance responds to lagged signed firm-specific returns as well as to market returns.

When the two factors are jointly incorporated into the model, the majority of individual stock volatilities are still found to be more asymmetric with respect to 'bad' systematic news than to 'bad' idiosyncratic news. Overall, the 'down-market' effect appears to be present in the volatility processes of the components of the DJIA and goes some way into explaining the observed difference in the asymmetry phenomenon between individual stocks and the parent index.

The remaining of the Chapter is organized as follows: Section 3.2 presents the data used in the empirical analysis, while Section 3.3 reports the results of estimating the standard GJR specification across the thirty sample stocks and the index. Section 3.4 discusses the modified GJR-I model in which market returns have replaced firm-specific ones, and Section 3.5 discusses the extended GJR-I model where the conditional variance responds to both market and stock returns. Finally, Section 3.6 concludes.

3.2 Data

The stock data used in this Chapter was obtained from the Wharton Research Data Services (WRDS) database. The data covers the decade from the 5th of January 1998 to the 31st of December 2007 for a total of 2,514 trading days. The sample consists of closing prices for the thirty components of the Dow Jones Industrial Average index across this ten-year period, based on the index's composition as of the 21st of November 2005. The thirty component stocks of the DJIA that are examined in this Chapter are presented in Table 3.1 which includes company name, company ticker, and date of entry in the index where applicable, i.e. if a firm entered the index after the beginning of the sample period.

Table 3.1

Composition of the DJIA

Company Name	Ticker	Date of Entry	Average Market Capitalization ($ billion)	Annualized Volatility
Aluminum Company of America (Alcoa)	AA		25.80	0.19
American International group Incorporated	AIG	08/04/2004	167.61	0.16
American Express Company	AXP		60.84	0.18
Boeing Company	BA		43.44	0.17
Citigroup	C		206.15	0.18
Caterpillar Incorporated	CAT		24.78	0.17
DuPont	DD		47.99	0.15
Walt Disney Company	DIS		52.90	0.18
General Electric Company	GE		371.16	0.15
General Motors Corporation	GM		24.01	0.20
Home Depot Incorporated	HD	01/11/1999	87.41	0.18
Honeywell International	HON	02/12/1999	31.90	0.19
Hewlett-Packard Company	HPQ		75.06	0.23
International Business Machines	IBM		158.45	0.17
Intel Corporation	INTC	01/11/1999	169.45	0.24
Johnson & Johnson Company	JNJ		161.69	0.12
J.P. Morgan & Company	JPM		88.05	0.19
Coca-Cola Company	KO		124.33	0.13
McDonald's Corporation	MCD		42.24	0.15
Minnesota Mining & Mfg (3M Company)	MMM		50.27	0.13
Philip Morris Companies Inc (Altria Group)	MO		113.56	0.16
Merck & Company Inc	MRK		119.98	0.15
Microsoft Corporation	MSFT	01/11/1999	318.26	0.18
Pfizer Incorporated	PFE	08/04/2004	194.10	0.16
Procter & Gamble Company	PG		142.33	0.13
AT&T Corporation	T		65.51	0.20
United Technologies Corporation	UTX		41.71	0.15
Verizon Communications Incorporated	VZ	08/04/2004	106.35	0.15
Wal-Mart Stores Incorporated	WMT		220.95	0.16
Exxon Corporation (Exxon Mobil)	XOM		304.20	0.13
Dow Jones Industrial Average	DJIA		3,640.46	0.09

This Table reports the constituent stocks of the Dow Jones Industrial Average as of the 21st of November 2005. The last column includes the (annualized) standard deviation of daily returns for each individual stock as well as for the parent index across the entire sample period January 1998 to December 2007.

The securities' closing prices in the WRDS files have been adjusted for dividends and for changes in the firms' capital structure. The GARCH specifications that are examined use daily log-returns across the thirty sample stocks, computed as the difference between the logarithms of two consecutive closing prices:

$$R_{i,t} \equiv \log(p_{i,t}) - \log(p_{i,t-1}) \tag{3.1}$$

where $R_{i,t}$ is the daily return of asset i at time t, and $p_{i,t}$ is the i^{th} asset's closing price at t. Daily closing prices for the Dow Jones index, adjusted for dividend distributions, were obtained from DataStream and corresponding logarithmic returns were computed using equation (3.1).

Table 3.2

Constituents Capitalization ($ million)

	Median	L_q	U_q	n
All Stocks	97,196	48,557	166,132	30
Survivors	75,062	45,711	150,390	23
Entrants	167,611	96,878	181,775	7

This Table tabulates summary statistics of the average market capitalisation of the thirty constituent stocks of the Dow Jones Industrial Average for the sample period January 1998 to December 2007. The market capitalisation of each stock is computed as the geometric mean of the year-end market capitalisation across the sample period. L_q and U_q refer to the first and to the third quartile, respectively, while n denotes the number of stocks included in the respective group. The terms 'Survivors' and 'Entrants' refer to companies that were included in the index's composition throughout the entire period and to companies that entered the index after the 5th of January 1998, respectively.

Table 3.2 presents summary statistics on the sample firms with respect to their size. Size is expressed in millions of dollars and it is proxied by the geometric mean of each firm's market capitalisation at year end across the ten-year sample period. As can be seen from the Table, the median firm has an average market capitalisation of $97.20 billion although size varies significantly across the thirty firms. A quarter of the stocks have market capitalisations less than $48.56 billion while another quarter of the stocks have an average market value of more than $166.13 billion. Furthermore, there is a obvious difference between the size of firms that were included in the index throughout the entire period, i.e. 'Survivors', and that of firms that entered the index after the beginning of the sample period, i.e. 'Entrants'. Firms in the 'Entrants' group are clearly larger than those in the 'Survivors' group, with the median 'entrant' having a market capitalisation of $167.61 billion, compared to $75.06 billion for the median 'survivor'.

Table 3.3 reports descriptive statistics on the returns distributions of the Dow Jones and of the thirty constituent stocks. The first four moments, the minimum and maximum values of the returns distribution are presented for the index as well as for the median, lower and upper quartile stocks, while the last column of the Table tabulates the number of stocks for which the respective statistic has a value greater than that of the index. Throughout the period January 1998 to December 2007, the DJIA appreciated by an average of 0.88 basis points per day, compared to a slightly higher average daily return of 0.92 bps for the median stock. Not surprisingly, as a result of diversification, the index was characterized by a much smaller variance of returns (0.22 bps on a daily basis) than all thirty stocks. The lower return variability in the case of the DJIA is also evidenced by the smaller (absolute) magnitudes of the maximum and minimum values, with all thirty individual stocks exhibiting higher absolute extreme observations. Finally, although index skewness is very close to that of the median stock, all thirty individual returns distributions are significantly more leptokurtotic (kurtosis for the lower quartile stock is 6.71) than that of the Dow Jones (kurtosis = 3.67).

Table 3.3

Summary Statistics of the Daily Returns on the DJIA and its Constituents

Summary Statistics	Index	Stocks			
		Median	L_q	U_q	Greater than Index
Mean x 10^4	0.88	0.92	0.31	1.27	17
Variance x 10^4	0.22	0.73	0.61	0.92	30
Skewness	-0.15	-0.14	-0.23	0.02	16
Kurtosis	3.67	7.68	6.71	10.43	30
Max	0.03	0.05	0.04	0.06	30
Min	-0.03	-0.06	-0.08	-0.05	0

This Table tabulates summary statistics of the daily returns of the DJIA index and of its thirty constituents. The sample period is 5 January 1998 to 31 December 2007.

3.3 Volatility Asymmetry with respect to Firm-specific Returns

In this study, the volatility process of the Dow Jones constituent stocks is modelled by the Glosten, Jagannathan and Runkle (1993) extension of the ARCH specification. The GJR(1,1) model is defined by the following equations for the mean and for the conditional variance of the returns generating process (see also equation (2.12) in the previous Chapter):

$$R_t = \mu + \sqrt{h_t} z_t \qquad z_t \sim i.i.d.N(0,1) \tag{3.2}$$

$$h_{i,t} = \alpha_0 + \alpha_1 R_{i,t-1}^2 + \alpha_2 s_{i,t-1} R_{i,t-1}^2 + \beta h_{i,t-1} \tag{3.3}$$

where $s_{i,t-1}$ is a dummy variable that takes the value of one if the lagged stock return is negative, and that of zero otherwise. The asymmetry property is incorporated into the conditional variance equation by the use of two separate coefficients for lagged positive and negative returns, namely α_1 and α_2, respectively, such that α_1 describes the impact of lagged positive returns on the conditional variance while α_2 describes the additional impact of lagged negative returns. A positive and significant α_2 suggests that the stock variance's response to a negative return is more pronounced compared to a positive one of similar absolute magnitude, with volatility increasing by more following a depreciation than an appreciation of the same size.

In addition to jointly examining the magnitude and significance of α_1 and α_2, the extent to which the conditional variance responds asymmetrically to past returns of different signs can be determined by the 'asymmetry ratio', estimated using equation (3.4) (see also equation (2.14) in Chapter 2). As has already been mentioned, lower values of this ratio suggest a more pronounced asymmetry in the conditional variance.

$$\text{Asymmetry Ratio} = \frac{\alpha_1}{\alpha_1 + \alpha_2} \tag{3.4}$$

Table 3.4 presents the results from estimating the standard GJR(1,1) model on the returns of the DJIA as well as on those of its thirty components in Panels A and B, respectively. In the case of the index, the intercept α_0 is significant with a p-value of 0.00, albeit small in magnitude, and the coefficient α_1 for lagged positive returns is statistically indistinguishable from zero (p-value is 0.72). Moreover, the estimated parameters indicate that the volatility process of the parent index is highly asymmetric since the coefficient α_2 for the additional impact of lagged negative returns is relatively high (0.12) and statistically significant (p-value is 0.00). This pronounced asymmetry in index returns can be further evidenced by the low 'asymmetry ratio' which takes the value of -

0.02. Finally, the coefficient β for the lagged conditional variance is 0.9318 and the volatility persistence of the process is estimated at 0.9892 ($\alpha_1 + \frac{\alpha_2}{2} + \beta$).

Table 3.4

Estimation of the Standard GJR

$$h_{i,t} = \alpha_0 + \alpha_1 R_{i,t-1}^2 + \alpha_2 s_{i,t-1} R_{i,t-1}^2 + \beta h_{i,t-1}$$

Panel A: Estimated Parameters for the Index

	α_0	α_1	α_2	β	Ratio
	0.0000	-0.0026	0.1200	0.9318	-0.02
	(0.00)	(0.72)	(0.00)	(0.00)	

Panel B: Estimated Parameters for the Constituent Stocks

	α_0	α_1	α_2	β	Ratio
Median Coefficient	0.0000	0.0517	0.0466	0.9072	0.23
Greater than Index	29	26	11	7	26

This Table reports the results from estimating the standard GJR(1,1) model on the returns of the DJIA and on those of its constituent stocks. The sample runs from January 1998 to December 2007. P-values on the coefficients of the index in Panel A are given in brackets. 'Ratio' is a measure of volatility asymmetry.

As can be seen from Panel B, the conditional variance processes of the thirty constituents of the Dow Jones differ significantly from that of the parent index, particularly with respect to the 'volatility asymmetry' effect. The median intercept is very close to zero but statistically significant, and 29 out of the 30 stocks have intercepts that are higher than that of the index. In contrast to the DJIA, the effect of past positive returns is statistically significant for the majority (26) of the stocks, while α_1 for the median stock is 0.0517. More importantly, though, individual stocks' variances appear to be less asymmetric compared to the index. The coefficient for the additional impact of lagged negative returns on the conditional variance, as measured by α_2, is 0.0466 for the median stock, compared to 0.1200 for the Dow Jones. Also, more than two thirds of the sample stocks have α_2s that are smaller than that of the index, suggesting that negative returns in the case of individual stocks have a less extreme incremental impact on volatility than that which is the case for the index.

The higher value of the 'asymmetry ratio' provides further support for the above finding. More specifically, the median 'asymmetry ratio' is 0.23, compared to -0.02 for the DJIA, while 26 out of the thirty stocks have higher (or less negative) 'ratios' than that of the index and, given that lower 'ratio' values are associated with higher asymmetry, it is concluded that volatility asymmetry across the vast majority of the DJIA's constituent stocks is significantly less pronounced than that of the parent index.

3.4 Volatility Asymmetry with respect to Market Returns

The previous Section documented that the volatility processes of almost all the individual components of the Dow Jones are characterized by a less severe asymmetry with respect to past returns of opposite signs when compared to the volatility process of the index. This Section examines the possibility of individual stocks' variances having an asymmetric response to 'news' on the market as opposed to firm-specific 'news'. A modified GARCH model is fitted on the returns of the thirty DJIA constituents, which, similarly to the standard GJR, includes signed lagged returns as regressors in the variance equation. However, contrary to the GJR, the above returns are market returns, proxied by the returns of the Dow Jones. The mean equation of this modification of the GJR, termed GJR-I, is the same as in the standard model, while the variance equation is given in (3.5):

$$h_{i,t} = \alpha_0 + \alpha_1^{IND} R_{IND,t-1}^2 + \alpha_2^{IND} s_{IND,t-1} R_{IND,t-1}^2 + \beta h_{i,t-1} \qquad (3.5)$$

where $R_{IND,t}$ is the return of the DJIA index at time t, and the dummy variable $s_{IND,t-1}$ takes the value of one if the lagged index return $R_{IND,t-1}$ is negative and the value of zero otherwise. Within this framework, the firm-specific volatility process responds to market shocks but not to firm-specific shocks, suggesting that stock volatility increases at down-markets and increases at up-markets, while past stock returns do not affect stock volatility. Similarly to the standard GJR specification, α_1^{IND} measures the symmetric impact of new systematic information, and α_2^{IND} captures the additional impact of 'bad' systematic information, i.e. of a negative lagged market return.

Panel A of Table 3.5 presents the results from estimating the GJR-I model across the thirty components of the Dow Jones. The Table reports the estimated coefficients of interest for the median stock as well as the number of stocks for which the respective parameter has a value that is greater than that of the index. Obviously, in the case of the DJIA the GJR-I is equivalent to the standard GJR specification since market returns are also asset-specific returns for the market index, so that the estimated GJR-I parameters for the index are the same as those reported in Panel A of Table 3.4. Moreover, the modified 'asymmetry ratio' is given by equation (3.6) where α_1^{IND} and α_2^{IND} have replaced α_1 and α_2, respectively.

$$\text{Modified Asymmetry Ratio} = \frac{\alpha_1^{IND}}{\alpha_1^{IND} + \alpha_2^{IND}} \tag{3.6}$$

The median intercept α_0 is again very close to zero but remains statistically significant for most of the sample stocks. The coefficient α_1^{IND} for the symmetric response of the conditional variance to lagged systematic innovations is 0.0157 for the median stock, which is relatively low when compared to its respective value of 0.0517 in the standard GJR specification. It appears, thus, that the symmetric impact of lagged index returns on the volatility of individual stocks is, on average, smaller than that of lagged firm-specific returns. However, the asymmetric effect is more pronounced with respect to market returns. The median α_2^{IND} is 0.1588, compared to a α_2 of 0.0466 for the median stock in the GJR (see Panel B of Table 3.4), indicating that the additional response of individual stock volatility to 'bad' market news is more than three times the magnitude of its response to 'bad' idiosyncratic news. Furthermore, the effect of negative lagged index returns for the median stock is higher even when compared to the index's conditional variance, for which the respective coefficient α_2 was 0.12, while more than half of the sample stocks (17 out of 30) have α_2^{IND} coefficients that are higher than the α_2 of the index.

The extent to which individual stock return volatilities respond asymmetrically to past index returns of opposite signs can be directly examined through the modified 'asymmetry ratio' which, as can be seen from Table 3.5, is 0.07 for the median stock. Given that the median standard asymmetry measure was 0.23, and that lower values are associated with a more pronounced volatility asymmetry, it appears that the conditional variance of the median stock is characterized by a more severe asymmetry with respect to negative market returns compared to negative firm-specific returns. In addition, when the modified measure is considered, only 17 out of the 30 individual stocks have higher ratios than that of the DJIA and, thus, exhibit lower asymmetry than that of the index. This result differs considerably from the one obtained when the standard GJR specification was estimated, when the majority of stock volatilities (26 out of 30) were found to be less asymmetric than the index's volatility.

Table 3.5

Estimation of the GJR-I

Panel A: Estimated Parameters for the GJR-I

$$h_{i,t} = \alpha_0 + \alpha_1^{IND} R_{IND,t-1}^2 + \alpha_2^{IND} s_{IND,t-1} R_{IND,t-1}^2 + \beta h_{i,t-1}$$

	α_0	a_1^{IND}	a_2^{IND}	Ratio
Median Coefficient	0.0000	0.0157	0.1588	0.07
Greater than Index	21	17	17	17

Panel B: Estimated Parameters for the Extended GJR-I

$$h_{i,t} = \alpha_0 + \alpha_1 R_{i,t-1}^2 + \alpha_2 s_{i,t-1} R_{i,t-1}^2 + \alpha_1^{IND} R_{IND,t-1}^2 + \alpha_2^{IND} s_{IND,t-1} R_{IND,t-1}^2 + \beta h_{i,t-1}$$

		Median	L_q	U_q	Greater than Index
	a_1	0.0497	0.0094	0.2140	25
Stock Returns	`a_2	-0.0030	-0.0871	0.0814	7
	Ratio	0.41	0.01	2.18	24
	a_1^{IND}	0.0329	-0.0364	0.1988	21
Index Returns	a_2^{IND}	0.1356	0.0202	0.2709	16
	Ratio	0.34	0.00	0.46	23

This Table reports the estimated parameters of the modified GJR-I specification across the thirty constituent stocks of the DJIA from January 1998 to December 2007. Panel A refers to the standard GJR-I, while Panel B refers to the extended GJR-I model. L_q and U_q denote the lower and upper quartiles, respectively. Ratio is a measure of volatility asymmetry.

3.5 Volatility Asymmetry with respect to both Market and Stock Returns

The previously examined GJR-I specification describes a stock's variance process under the relatively limiting assumption that volatility responds only to market movements and that it is uncorrelated with firm-specific stock movements. In order to relax this limitation, an extension of the GJR-I model is estimated, where the individual stock's conditional variance responds asymmetrically to past index returns as well as to past firm-specific returns. The equation of the conditional variance in the extended GJR-I is given in (3.7) and the estimated parameters of the new specification across the thirty DJIA components are presented in Panel B of Table 3.5. The Table reports the median, lower and upper quartile values of the coefficients of interest, i.e. the coefficients that capture the impact of past signed firm-specific and index returns, across the thirty sample stocks. The last column of Panel B tabulates the number of stocks for which the respective coefficient has a value that is greater than that of the index. The results are grouped into two sets, where the first set refers to firm-specific innovations and includes α_1, α_2 and the standard asymmetry measure, while the second one refers to market innovations and includes α_1^{IND}, α_2^{IND} and the modified asymmetry measure.

$$h_{i,t} = \alpha_0 + \alpha_1 R_{i,t-1}^2 + \alpha_2 s_{t-1} R_{i,t-1}^2 + \alpha_1^{IND} R_{IND,t-1}^2 + \alpha_2^{IND} s_{IND,t-1} R_{IND,t-1}^2 + \beta h_{i,t-1} \qquad (3.7)$$

When the effects of idiosyncratic and of systematic innovations are jointly examined, the difference in the volatility's asymmetric response to the two factors somewhat decreases. The coefficient α_1 that captures the impact of lagged positive firm-specific returns is found to be generally higher than the respective coefficient α_1^{IND} that captures the impact of lagged positive market returns. More specifically, the median α_1 is 0.0497, compared to a median α_1^{IND} of 0.0329, while similarly the lower and upper quartiles of α_1 are both higher than those of α_1^{IND} (0.0094 and 0.2140 compared to -0.0364 and 0.1988, respectively).

On the other hand, the additional impact of lagged negative returns is clearly still higher in index returns than in individual stock returns. The median α_2^{IND} is 0.1356 while the median α_2 is very close to zero and slightly negative at -0.0030. Furthermore, the asymmetry 'ratio' is on average higher, i.e. less asymmetric, than the respective modified asymmetry measure. The median stock has a ratio of 0.41, compared to a modified ratio of 0.34, while 19 out of the 30 sample stocks exhibit a more pronounced asymmetry with respect to past index returns than to past firm-specific returns.

Overall, the thirty components of the Dow Jones are found to be characterized on average by a more pronounced volatility asymmetry with respect to 'bad' market news than to 'bad' idiosyncratic news. In other words, negative market innovations appear to have a more extreme additional impact on the conditional variance of individual stock returns compared to negative firm-specific innovations. Also, when volatility asymmetry is examined within the GJR-I framework of market returns, the results cast some doubt on the previously reported empirical finding of stock volatilities being less asymmetric than those of indices.

3.6 Conclusion

This Chapter has examined the volatility processes of the individual stocks that are included in the Dow Jones index, from the perspective of the asymmetric response of volatility to past returns of opposite signs. Previous studies have documented that individual stock variances are usually less asymmetric compared to index variance and, although little has been done towards explaining this difference, this has been generally accepted as a property of the returns generating process of individual stocks.

When the standard GJR model is fitted on the returns of the thirty DJIA constituents, the results are in line with previous empirical findings in the sense that, even though volatility asymmetry varies across the sample stocks, the vast majority of the components of the Dow Jones are less asymmetric than the index. In addition to a moderate volatility asymmetry with respect to past firm-specific returns, this study proposes that individual stocks' conditional variances respond asymmetrically to signed past returns of the market. Estimating a modification of the GJR model, termed GJR-I, where signed lagged index returns have replaced firm-specific returns, indicates that stock returns are significantly correlated with market returns. Furthermore, individual stock volatility is found to increase by more following 'bad' market news compared to 'good' market news, and this asymmetry is on average more pronounced compared to the one indicated by the standard GJR.

The above results still hold after estimating an extension of the GJR-I specification, where the conditional variance of individual stocks is a function of signed lagged firm-specific returns as well as of market returns. When the two factors are jointly incorporated into the model, the thirty individual stock volatilities are generally correlated with both idiosyncratic and systematic innovations, but the additional impact of negative systematic innovations on the conditional variance is on average higher than the additional impact of negative idiosyncratic innovations. Overall, the empirical findings of this Chapter are consistent with the hypothesis that volatility asymmetry is in essence a 'down-market effect', not necessarily attributed to changes in leverage or to volatility feedback.

References

Bekaert, Geert, and Guojun Wu, 2000, Asymmetric Volatility and Risk in Equity Markets, *Review of Financial Studies* 13, 1-42

Black, Fisher, 1976, Studies of Stock Price Volatility Changes, *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section*, 177-181

Blair, Bevan, Ser-Huang Poon, and Stephen Taylor, 2002, Asymmetric and Crash Effects in Stock Volatility for the S&P 100 Index and its Constituents, *Applied Financial Economics* 12, 319-329

Bollerslev, Tim, Julia Litvinova, and George Tauchen, 2006, Leverage and Volatility Feedback Effects in High-Frequency Data, *Journal of Financial Econometrics* 4, 353-384

Campbell, John, and Ludger Hentschel, 1992, No News is Good News, *Journal of Financial Economics* 31, 281-318

Christie, Andrew, 1982, The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects, *Journal of Financial Economics* 3, 407-432

Figlewski, Stephen, and Xiaozu Wang, 2000, Is the 'Leverage Effect' a Leverage Effect?, *NYU working paper*

French, Kenneth, G. William Schwert, and Robert Stambaugh, 1987, Expected Stock Returns and Volatility, *Journal of Financial Economics* 19, 3-30

Kim, Dogncheol, and Stanley Kon, 1994, Alternative Models for the Conditional Heteroscedasticity of Stock Returns, *Journal of Business* 67, 563-598

Smith, Daniel, 2007, Risk and Return in Stochastic Volatility Models: Volatility Feedback Matters!, *Simon Fraser University working paper*

Stivers, Christopher, Patrick Dennis, and Stewart Mayhew, 2006, Stock Returns, Implied Volatility Innovations and the Implied Volatility Phenomenon, *Journal of Financial and Quantitative Analysis* 41, 381-406

Tauchen, George, Harold Zhang, and Ming Liu, 1996, Volume, Volatility and Leverage: A Dynamic Analysis, *Journal of Econometrics* 74, 177-208

Chapter 4

An Examination of the Forward Premium Puzzle Using Option Implied Information

4.1 Introduction

4.1.1 Literature Review

The forward premium puzzle refers to the widely observed rejection of the forward premium as a conditionally unbiased predictor of future spot exchange rate returns. When exchange rate returns are regressed on the lagged forward premium/discount, interest rate parity predicts a slope coefficient equal to one. However, a vast body of the related literature reports a coefficient less than the theoretical value of unity and, in most cases, significantly negative. For instance, McCallum (1994), using yen, mark and pound rates against the dollar for the period 1978-1990, estimates a slope coefficient of -4, an estimate that is considered typical of such studies. Other studies that report negative coefficients include Backus, Gregory and Telmer (1993), Mark, Wu and Hai (1993), Froot and Frankel (1989), and Byers and Peel (1991). According to the *Uncovered Interest Rate Parity* (UIP) condition, this translates to an appreciating currency for the country with the higher nominal interest rate. Finally, Hodrick (1987) and Engel (1996) provide comprehensive surveys of the forward premium puzzle literature.

The variety of explanations that have been suggested to account for the forward premium anomaly indicates that researchers so far have failed to reach a consensus with respect to the source of this negative correlation between the forward rate and exchange rate returns. One of the most common explanations proposed has been the presence of a time-varying risk premium, defined as the difference between the forward rate and the conditional expectation of the future spot rate, with the risk premium assumed to be negatively correlated with expected spot rate returns. More specifically, one stream of the literature has taken the risk premium to be exogenous and attempted to explore whether its variance and magnitude can account for some of the forward premium bias. Boyer and

Adams (1988) and McCallum (1994) have focused on this direction and, although their models were able to produce a negative value of the forward premium's slope coefficient, Engel (1996) argues that '...it seems unlikely that the behaviour of rp_t^{re} (risk-premium at time t under rational expectations) in these models could be reconciled with existing models of risk-averse behaviour'.

Another line of thought has focused on the presence of a systematic forecast error in the forward premium that is related to the way investors form expectations about future levels of exchange rates. For instance, Lewis (1994) and Evans (1995) argue that previous empirical findings might be inconclusive due to the use of small samples and the occurrence of infrequent extreme events, a combination that is likely to make spot returns appear predictable (a phenomenon that is usually referred to as the 'peso problem'). Moreover, Gourinchas and Tornell (2004) develop a model where markets determine nominal exchange rates according to whether interest rate shocks are perceived to be transitory or persistent. Within this framework, Gourinchas and Tornell (2004) suggest that a large portion of the forward rate bias can be explained if '...investors misperceive shocks to be more transitory than what they actually are'. Finally, Bacchetta and Wincoop (2007, 2008) examine the forward premium puzzle from the perspective of managing currency portfolios. Assuming that active management of short-term currency positions can be problematic due to the 'unpredictability' of future spot rates, Bacchetta and Wincoop find that infrequent rebalancing of FX portfolios combined with incomplete information processing can lead to 'delayed overshooting', where exchange rates continue to adjust to changes in interest rates long after the initial shock.

In contrast to the above negative bias being driven by investors' expectation-forming process, other papers suggest monetary policy interventions as a potential factor. More specifically, assuming that monetary authorities set a target band of exchange rates, they are likely to intervene after a large shock in the spot rate by changing the interest rate in order to offset the shock and return to the target band (see McCallum (1994)).

A number of more recent studies has documented that the predictive ability of the forward rate on future spot rates appears to be significantly related to the forecasting time-horizon. This line of research suggests that, although the anomaly is pronounced when medium-term forward rates are quoted (one month up to one year), this is not necessarily the case for shorter or longer maturities. With respect to short-term contracts, Chaboud and Wright (2005) examine a set of high-frequency data on five currencies vis-à-vis the US dollar, and find that the slope coefficient of future spot rates regressed on forward rates is close to its theoretical value of unity for short windows of up to one day. In addition, Bernoth, Hagen and Vries (2007) use data from the futures market to show that the correlation between the forward premium and realized exchange rate returns starts positive and relatively close to one for the shortest maturities, and slowly turns negative as maturity approaches the monthly level. At the other end of the spectrum, Chinn and Meredith (2004) focus on multi-year forward rates of six currencies with respect to the US dollar. They report that, despite the failure of interest parity in the short run, the slope coefficient is significantly positive and closer to unity than to zero over longer horizons, possibly as a result of exchange rates being driven by fundamentals rather than short-term speculation (see also Alexius (2001) and Chinn (2006)).

In order to examine the predictive power of the forward rate, some of the earlier studies estimated a simple model where the log of the future spot rate is regressed on the log of the forward rate through Ordinary Least Squares (OLS) minimization. Subsequent research, though, on the time-series properties of the above variables has demonstrated that this model is potentially misspecified. In one of the most influential papers, Baillie and Bollerslev (1994) examine a set of seven currencies and their respective thirty-day forward rates, and find that there is strong evidence of cointegration between spot and forward rates. Given, thus, that the stationarity assumption of the OLS framework is potentially violated in spot-forward regressions, it has been argued that statistical inference with respect to the resulting slopes is relatively problematic.

After the above research in unit roots, most of the surveys examining the forward premium anomaly have explored the typical regression model in which exchange rate returns are the dependent variable and the lagged forward premium/discount is the only explanatory variable. However, it has been noted that under the moderate assumption of log-normal distribution for exchange rates, forward prices and price levels, two correction terms must be added to the regression specification. Those terms are related to the variance of the spot rate and its covariance with the price level, and are commonly referred to as the *Jensen's Inequality Terms* (JIT). While this correction is dictated by theory, its empirical effect has been frequently questioned. For instance, Bekaert and Hodrick (1993) report that including the variance correction term does not result in a slope coefficient that is consistent with theoretical predictions, while similar conclusions have been reached by Cumby (1988), Hodrick (1989b), Baillie and Bollerslev (1990), and Backus, Gregory and Telmer (1993).

4.1.2 Scope of Study

This Chapter examines the forward premium puzzle with particular emphasis on the role of option-implied information in explaining this widely documented anomaly. Given that previous regressions of future spot rates (returns) on forward rates (forward premium) have resulted in estimated slopes that are significantly different from the value of unity that UIP predicts, and even become negative in certain cases, the hypothesis of the omitted future spot variance at least partly driving the results is explored.

Assuming lognormally distributed exchange rates, Jensen's Inequality results in a time-varying risk-premium in the forward markets that incorporates a term that refers to the future variance of exchange rate returns. The above correction has received relatively little attention in the existing literature, with a number of researchers arguing that its overall effect in accounting for deviations from UIP is not significant. The motivation for this study stems from the fact that previous papers have used historical measures of volatility to proxy for the spot rate's variance. However, since the variable in the extended regression specification refers to future volatility, it could be the case that the observed failure of JIT in estimating a forward slope closer to its theoretical value of one might be due to the use of a poor proxy for future variance. The subsequent analysis attempts to correct this potential limitation by estimating forward-looking volatility proxies, namely volatilities implied by currency options prices, which have been shown to have significant forecasting power over future volatility in foreign exchange markets (see for instance Pong, Shackleton, Taylor and Xu (2004)).

Overall, the results seem to indicate that including the spot rate's future variance in an extended specification provides forward slopes that are significantly closer to one compared to the standard univariate model. More specifically, incorporating the implied spot variance as an additional regressor and using the Fully Modified Least Absolute Deviations (FM-LAD) estimator results in a three-fold increase in the proportion of beta estimates that fail to reject the null hypothesis of unity.

The remaining of the Chapter is organized as follows: Section 4.2 gives an overview of the economic relationships leading to the Uncovered Interest Parity condition that is empirically tested. Section 4.3 describes the data used and the time-series properties of the pound/dollar exchange rate and its monthly forward rate. Section 4.4 presents the three alternative methodologies used to extract the spot rate's expected variance from option prices, while Section 4.5 describes the results from the Ordinary Least Squares estimation of the forward unbiasedness hypothesis. Section 4.6 examines the limitations of estimating the specification using OLS and discusses the FM-LAD technique and its results. Finally, Section 4.7 concludes.

4.2 Economics

Throughout this Chapter, S_t denotes the spot exchange rate at time t, while $F_t^{t+\tau}$ refers to the forward exchange rate at time t for delivery at time $t+\tau$. Corresponding logarithmic values are denoted by the lower case variables s_t and $f_t^{t+\tau}$, respectively. Both rates use the US dollar as the numeraire currency. Furthermore, i_t refers to the risk-free interest rate applicable for US investors, while i_t^* denotes the foreign risk-free rate.

The *Covered Interest Rate Parity* condition (CIP) states that the difference between the forward rate and the spot rate at time t must be equal to the interest rate differential between the two countries. There is strong empirical evidence demonstrating that, ignoring transaction costs, CIP generally holds (see for instance Bahmani-Oskooee and Das (1985) and Clinton (1988)).

$$f_t^{t+\tau} - s_t = i_t^* - i_t \tag{4.1}$$

The *Uncovered Interest Rate Parity* condition (UIP) then states that the expectation of spot rate returns must be equal to the interest rate differential. Taking into account (4.1), UIP can be expressed as follows,

$$E_t[s_{t+\tau} - s_t] = f_t^{t+\tau} - s_t = i_t^* - i_t \tag{4.2}$$

where $E_t(\cdot)$ is a (risk-neutral) expectation operator conditional on information available at time t. This study focuses on the prediction that the expected future spot rate $E[s_{t+\tau}]$ must be equal to the current forward rate $f_t^{t+\tau}$, or equivalently, that expected spot returns $E_t[\Delta s_{t+\tau}]$ must be equal to the current forward premium $f_t^{t+\tau} - s_t$.

In order to derive UIP, one must jointly assume rational risk-neutral agents, free capital mobility and the absence of taxes on capital transfers. From the same set of assumptions, it is also implied that expected real returns from trading in the forward market must be zero.

$$E_t[\frac{F_t^{t+\tau} - S_{t+\tau}}{P_{t+\tau}}] = 0 \tag{4.3}$$

where $P_{t+\tau}$ denotes the domestic dollar price level at time $t+\tau$. Assuming that all three variables in (4.3) are lognormally distributed and by using a Taylor series expansion to second terms, equation (4.4) is derived

$$E_t[s_{t+\tau}] - f_t^{t+\tau} = -\frac{1}{2}\mathrm{var}_t(s_{t+\tau}) + \mathrm{cov}_t(s_{t+\tau}, p_{t+\tau}) \tag{4.4}$$

where $p_{t+\tau}$ is the logarithm of the price level $P_{t+\tau}$. The above two conditional second moment terms are usually referred to as *Jensen's Inequality Terms* (JIT). A more detailed discussion on the derivation of (4.4) under a *Stochastic Discount Factor* framework is provided in Appendix A of this Chapter, while similar derivations are also given in Azar (2008), Engel (1999), and Soderlind and Svensson (1997). This extended specification suggests that the time-varying risk-premium in foreign exchange markets that has been examined in previous papers, i.e. the difference between the forward rate and the subsequent realization of the spot rate, depends on the variance of the spot rate as well as on the covariance between the spot rate and the domestic price level. In contrast to equity markets where a risk-free asset (cash) is exchanged with risky assets (stocks), though, currency markets are based on investors 'swapping' risky assets. Therefore, it is not straightforward to determine whether the domestic or the foreign investor will receive this foreign exchange risk-premium at a given point in time as this will depend on macroeconomic fundamentals. Finally, it should be noted that the above *Jensen's Inequality Terms* are shown to be directly related to the foreign exchange risk-premium not necessarily in terms of a rigorous theoretical framework but rather as a result of a mathematical paradox, i.e. *Siegel's Paradox*, which is based on the convexity property of exchange rates as ratios and on the concavity property of the logarithmic function. Equation (4.4) can then be rewritten as

$$E_t[s_{t+\tau}] = f_t^{t+\tau} - \frac{1}{2}\mathrm{var}_t(s_{t+\tau}) + \mathrm{cov}_t(s_{t+\tau}, p_{t+\tau}) \tag{4.5}$$

or, in terms of returns,

$$E_t[s_{t+\tau}] - s_t = (f_t^{t+\tau} - s_t) - \frac{1}{2}\mathrm{var}_t(s_{t+\tau}) + \mathrm{cov}_t(s_{t+\tau}, p_{t+\tau}) \tag{4.6}$$

which describe the models that have been examined in past surveys. In this Chapter, the emphasis is on the effect of future variance $var_t(s_{t+\tau})$ in explaining future spot levels, not taking into account the covariance between spot rates and the price level. Although theory predicts that the latter variable will have some explanatory power in predicting future exchange rates, it has been argued that its measurement is relatively problematic, reducing its actual explanatory power. More specifically, price levels, like many other economic variables, are reported at relatively low frequencies. The resulting smoothing and averaging complicates its inclusion in the regression model, especially considering the fact that the dependent variable as well as the remaining explanatory variables are estimated at a daily frequency.

Also note that the main variable of interest, namely $var_t(s_{t+\tau})$, refers to the future period $t+\tau$ and is, therefore, not observable at time t. One methodology that has been used in the related literature involves fitting an historical model to past data and inferring future volatility through this model's parameters. However, since the JIT measures the variance at $t+\tau$, a forward-looking measure, such as implied variance, might be a more appropriate proxy for future variance. In order to obtain a forward-looking measure of $var_t(s_{t+\tau})$, implied variances are extracted from a set of options written on foreign exchange using three different techniques. The methodology for estimating implied volatility from option prices is described in detail in Section 4.4.

4.3 Data

4.3.1 Sources

This study focuses on the exchange rate of the British pound vis-à-vis the US dollar. The sample period runs from January 1993 to March 2000, for a total of 1,833 trading days. Daily spot exchange rates and 30-day forward rates at a daily frequency (proxied by exchange-traded futures rates) were obtained from DataStream.

The original options dataset comprised of a total of 157,733 options written on the pound/dollar exchange rate, with the dataset containing, among other fields, option prices, strike prices, time-to-maturity, implied volatilities and trading volume. Prices of foreign exchange options are calculated as the mid-point of the best bid and the best ask quote at the end of the trading day, while option implied volatilities are calculated using the Black and Scholes (1973) option pricing formula.

Similarly to previous studies, several filters were introduced. First, all options with prices lying close to zero or outside the theoretical bounds were removed from the sample. Second, options that expired within a trading week (five trading days) were removed. Finally, observations with less than five traded contracts were also dropped to avoid illiquidity concerns. The above filtering resulted in a reduced dataset, with the final sample comprising of 44,645 options (22,939 calls and 21,706 puts).

The risk-free rate of interest is proxied by the LIBOR offered to US investors, whilst the 'dividend yield' of the underlying asset, i.e. of the spot exchange rate, is proxied by the UK LIBOR. A significant part of the related literature suggests that this is a reasonable proxy of the 'dividend yield' for an investor buying a currency option which gives her the right to buy the British pound using US dollars. The main intuition behind this choice is the fact that, had she instead bought the underlying, she would have been able to receive a return equal to the UK risk-free rate by investing in UK government bonds (see also Garman and Kohlhagen (1983), Chesney and Jeanblanc (2003), Boyrie, Kim and Pak (2005), and Cheng, Gallant, Ji and Lee (2005)) . However, it has to be noted that, within the context of this study, the term 'dividend yield' does not refer to an actual cash dividend paid by the underlying since this is not applicable to foreign currency. Instead, the foreign risk-free rate is essentially the equivalent of the dividend yield when pricing options in the sense that the investor holding an option contract rather than the underlying foregoes this risk-free payment. The US and UK LIBOR rates were obtained through DataStream.

4.3.2 Time-Series Properties

Figures 4.1 and 4.2 present the time-evolution of the pound/dollar spot exchange rate and forward rate, respectively, for the period January 1993 to March 2000. The spot rate ranged from a minimum of 1.42 to a maximum of 1.71, with the British Pound experiencing a relative appreciation with respect to the US Dollar during the overall sample period. Although the time-series of the monthly forward rate follows a similar pattern to s_t, it is also characterized by a number of extreme observations, especially during the 8-month period from July 1997 to March 1998. These outliers tend to be associated with a significantly positive forward premium $(f_t^{t+\tau} - s_t)$, which systematically fails to materialize as a higher future spot rate. Consequently, despite the fact that the minimum $f_t^{t+\tau}$ is very close to the minimum s_t, the range of the forward rate's level is significantly larger, with a maximum of 2.05.

Figure 4.1

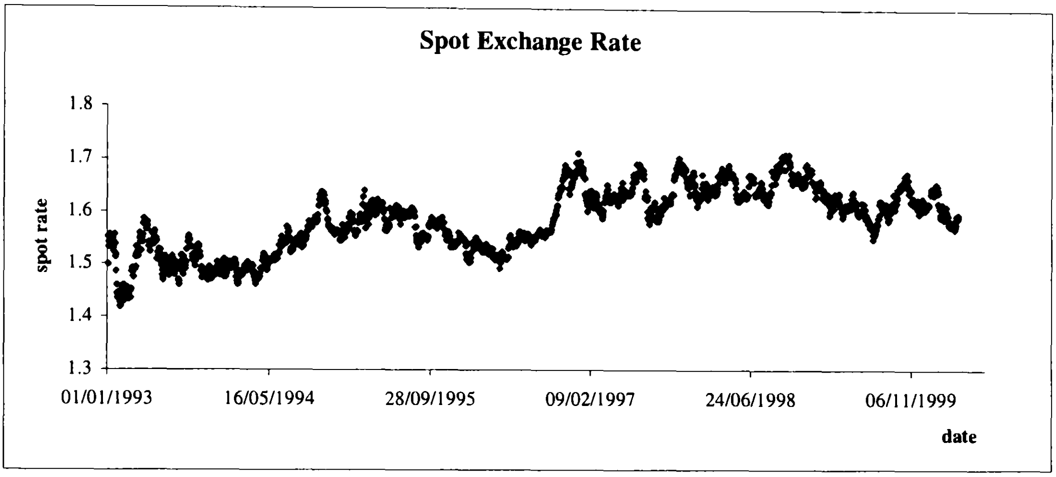

Figure 4.2

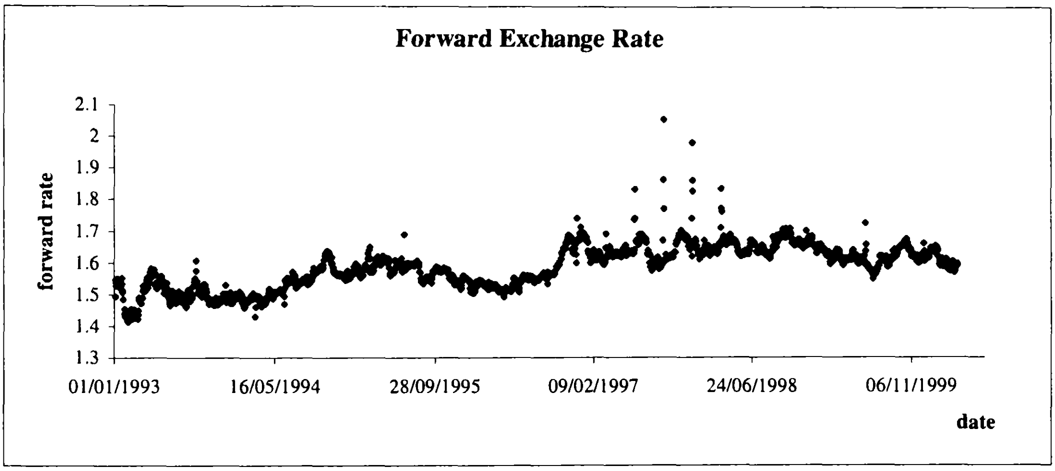

One consideration when examining the relationship between spot and forward rates is their respective orders of integration. A large part of the previous literature has focused on this property, with results not always in the same direction. This Section discusses the stationarity hypothesis of the main dependant and explanatory variables, namely the logarithmic spot exchange rate and forward rate, as well as of their first differences.

Many studies have examined the time-series properties of foreign currency exchange rates, finding s_t to follow a unit-root process, making foreign exchange returns

$(s_{t+\tau} - s_t)$ an I(0) process as the first difference of a unit-root. However, results on the order of integration of the forward rate $f_t^{t+\tau}$ and of the forward premium $(f_t^{t+\tau} - s_t)$ have been less than conclusive. Mark, Yu and Hai (1993) support the stationarity of the forward premium in their empirical investigation of three main exchange rates, namely the pound/dollar, French franc/dollar and yen/dollar rates. On the other hand, Crowder (1994) contradicts these results. Examining monthly observations for the pound, mark and Canadian dollar relative to the US dollar from January 1974 to December 1991, he finds that the null hypothesis of non-stationarity of $(f_t^{t+\tau} - s_t)$ cannot be rejected.

Figure 4.3

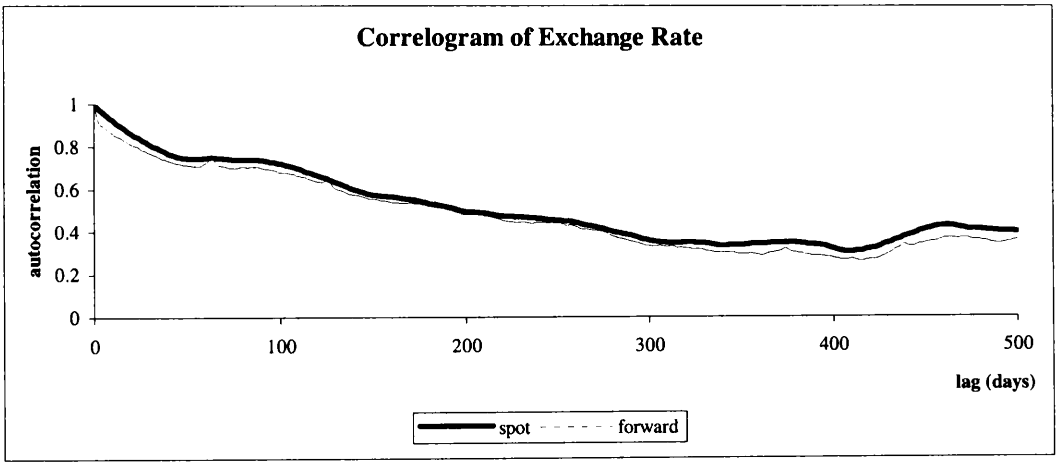

Within the sample period, monthly forward and spot rates exhibit very high, positive, slow-decaying autocorrelations in the first 100 lags. Moreover, the correlogram in Figure 4.3 indicates that autocorrelations for $f_t^{t+\tau}$ are at a slightly lower level than those of the spot rate. By subtracting s_t, exchange rate returns have significantly lower serial correlations after the first lags. However, autocorrelations in the forward premium $(f_t^{t+\tau} - s_t)$ are significantly higher than those of $(s_{t+\tau} - s_t)$, suggesting that the forward premium's deviation from stationarity is more pronounced.

Figure 4.4

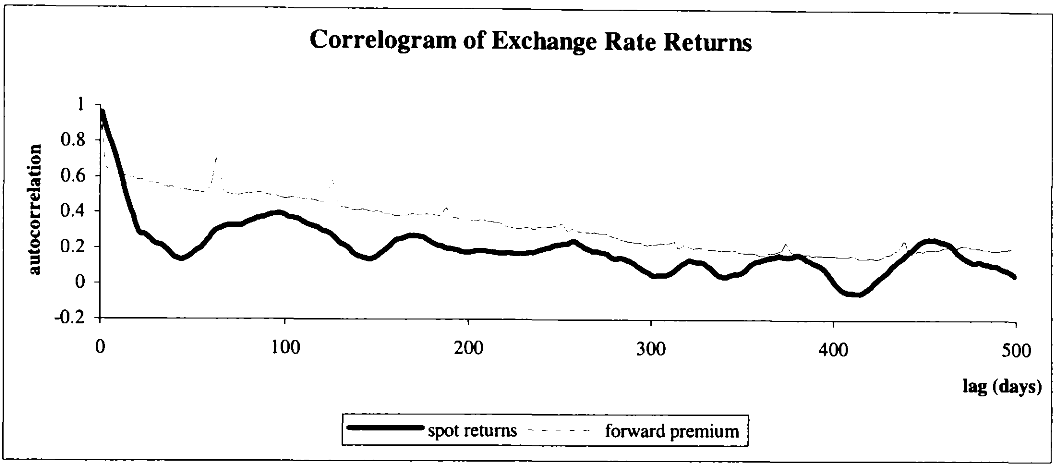

In addition to the correlograms, the order of integration of the variables of interest is examined by performing the Augmented Dickey-Fuller (ADF) test, the Kwiatkowski, Phillips, Schmidt, and Shin (KPSS) test, and the Geweke and Porter-Hudak (GPH) test, with the results presented in Table 4.1. The ADF tests the null that the series examined is a unit root I(1) and the number of lags is selected by the Schwartz Information Criterion. As can be seen from Panel A, the results support the presence of a unit root in the logarithmic spot exchange rate (t-stat = -2.76), but the null of a unit root in the forward rate process is rejected (t-stat = -3.19). Also, the hypothesis of the variables' first differences, i.e. $s_{t+\tau} - s_t$ and $f_t^{t+\tau} - s_t$, being unit roots is clearly rejected at the 5% level. Panel B suggests that the results of the KPSS test are consistent with those of ADF. More specifically, the KPSS tests the null that the series examined is stationary, with the number of lags being selected by Newey-West bandwidth using Bartlett kernel spectral levels. Spot rate returns $s_{t+\tau} - s_t$ are found to be stationary at the 5% level (t-stat = 0.06), which is to be expected given that they are the first difference of the unit root process s_t, while the null of stationarity is rejected for forward premium $f_t^{t+\tau} - s_t$ (t-stat = 0.77). Finally, the GPH test estimates that the order of integration of the spot rate is very close to unity (d = 1.04) but that the respective coefficient of the forward rate is significantly further from one (d = 0.89). The GPH results are also consistent with ADF and KPSS with respect to the series' first differences since spot rate returns $s_{t+\tau} - s_t$ are shown to be stationary (d = 0.00) but the forward premium $f_t^{t+\tau} - s_t$ might be characterized by a fractionally integrated process I(d), with d equal to 0.13.

Table 4.1

Stationarity Tests

	s_t	$f_t^{t+\tau}$	$s_{t+\tau} - s_t$	$f_t^{t+\tau} - s_t$
		Panel A: ADF test		
t-stat	-2.76*	-3.19	-7.93	-19.53
p-value	0.06	0.02	0.00	0.00
R-square	0.01	0.13	0.03	0.31
		Panel B: KPSS test		
t-stat	3.57	3.62	0.06*	0.77
Bandwidth	33	33	31	10
		Panel C: GPH test		
d	1.04	0.89	0.00	0.13

This Table tabulates the results of tests for the order of integration of the main (logarithmic) variables, namely the spot exchange rate, the forward rate, the spot rate returns and the forward premium. Panel A describes the results of the Augmented Dickey-Fuller (ADF) test, where the null is that the series examined is a unit root. Panel B describes the results of the Kwiatkowski, Phillips, Schmidt, and Shin (KPSS) test, where the null is that the series examined is stationary. Panel C reports the variables' order of integration (d) based on the Geweke and Porter-Hudak (GPH) test. The number of lags in the ADF test is obtained by the Schwartz Information Criterion. The number of lags in the KPSS test is selected automatically by Newey-West bandwidth using Bartlett kernel spectral levels. The 5% critical values are -2.86 and 0.46 for the ADF and KPSS tests, respectively. * denotes statistical significance at the 5% level.

Before proceeding with examining the forward unbiasedness hypothesis, attention should be drawn to the issue of potential data contamination in the time-series of spot and forward rates. This concern has been extensively discussed by Maynard and Phillips (2001), who report that forward premia series that are obtained from five alternative sources exhibit significant differences, tending to '... obfuscate the true time series properties of the forward premium, creating a clear (finite sample) bias in favour of stationarity'. Furthermore, Maynard and Phillips (2001) observe that the forward premium in their sample is characterized by large one-day fluctuations that are not present in the interest rate differential. This pattern is also found in the data used in this study as evidenced by the time-series of the nominal interest differential between the US and the UK, and of the forward premium, plotted in Figures 4.5 and 4.6, respectively. This implies a significant deviation from CIP, which predicts that the forward premium should be identical to the nominal interest rate differential.

Figure 4.5

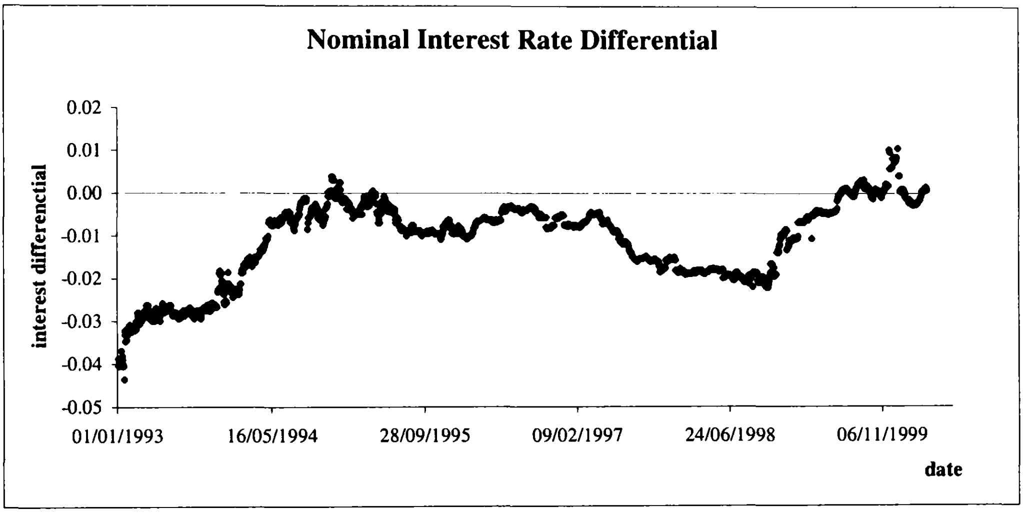

Figure 4.6

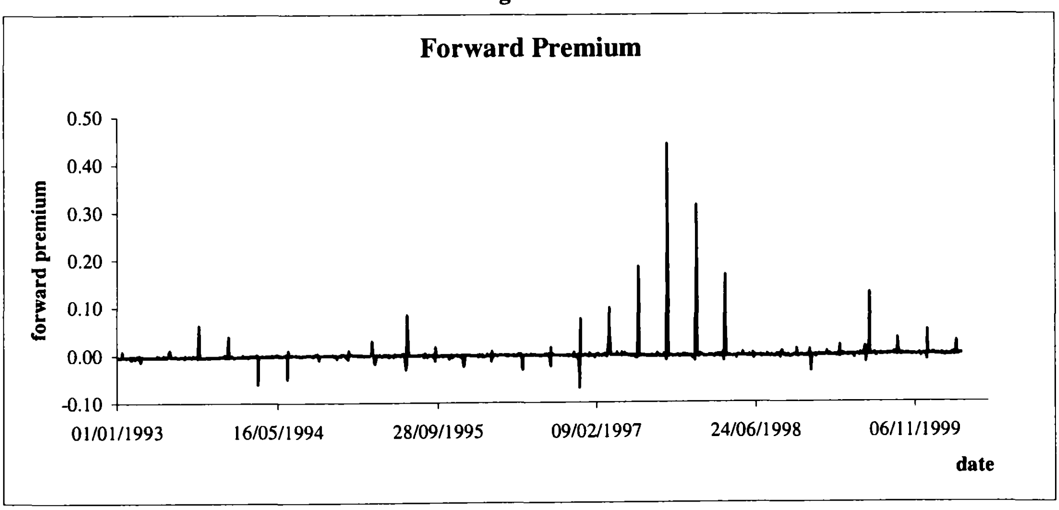

4.4 Estimating Option-Implied Variances

This Section describes the methodology used to estimate forward-looking measures of the spot rate's expected variance using option-implied information. Three different approaches are used to extract option-implied variances, namely the Black and Scholes (1973) option pricing framework, the Corrado and Su (1996) correction for non-normality, and the Britten-Jones and Neuberger (2000) model-free methodology.

4.4.1 Black-Scholes ATM Implied Variance

Assuming lognormally distributed asset returns, a constant risk-free rate and constant dividend yield, the Black and Scholes (1973) formula describes the price Q_{BS} of a European-style option as a function of the underlying asset's current price A_0, the option's exercise price K, the volatility σ of asset returns until the option's maturity, the risk-free rate r, the underlying's constant dividend yield q, and the time until the option's expiration T.

$$Q_{BS} = f(A_0, K, \sigma, r, q, T) \tag{4.7}$$

Within the Black and Scholes (B&S) framework, the price of a call C_{BS} and a put P_{BS} written on A are described by equations (4.8) and (4.9), respectively

$$C_{BS} = A_0 e^{-qT} N(d_1) - K e^{-rT} N(d_2) \tag{4.8}$$

$$P_{BS} = K e^{-rT} N(-d_2) - A_0 e^{-qT} N(-d_1) \tag{4.9}$$

where $N(\cdot)$ is the standard normal cumulative distribution function, and d_1 and d_2 are given by

$$d_1 = \frac{\ln(\frac{A_0}{K}) + (r - q + \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}$$

$$d_2 = d_1 - \sigma\sqrt{T}$$

The B&S formula has been widely used in the literature to infer the implied volatility $\hat{\sigma}$ of the underlying asset until the option's expiration. More specifically, given that the remaining five parameters A_0, K, r, q, T are readily observable and that market prices of options are available, the implied volatility of an option Q can be easily estimated by substituting theoretical prices Q_{BS} in the B&S formulas with quoted market prices Q_M of options, and then solving for σ. The resulting estimate $\hat{\sigma}$ then constitutes a forward-looking measure of the underlying's return volatility until the option's maturity T.

As has already been mentioned, the options dataset includes implied volatility curves of the spot rate across option strikes. First, option moneyness is expressed in terms of option delta, rather than of exercise price. More specifically, the delta Δ_i of option i is a measure of the option's local sensitivity to movements in the level of the underlying and is given as the first derivative of the option's price Q_i with respect to the underlying price A.

$$\Delta_i = \frac{\partial Q_i}{\partial A} \tag{4.10}$$

The analytical solution for the option delta under the B&S assumptions is described in equations (4.11) and (4.12), with Δ_C and Δ_P denoting the B&S deltas of calls and puts, respectively.

$$\Delta_C = e^{-qT} N(d_1) \tag{4.11}$$

$$\Delta_P = -e^{-qT} N(-d_1) \tag{4.12}$$

The first measure of implied volatility in this study is the Black and Scholes at-the-money (ATM) volatility. For each sample day, the first out-of-the-money (OTM) call and the first OTM put for the nearest expiration date T are identified. Then, using a linear interpolation across moneyness, the implied volatility σ_{ATM}^{T} of a hypothetical ATM option is estimated through the following equation:

$$\sigma_{ATM}^{T} = (\frac{\Delta_{p'}^{T} - 0.5e^{-qT}}{\Delta_{p'}^{T} - \Delta_{c}^{T}})\sigma_{OTM,c}^{T} + (\frac{0.5e^{-qT} - \Delta_{c}^{T}}{\Delta_{p'}^{T} - \Delta_{c}^{T}})\sigma_{OTM,p}^{T} \tag{4.13}$$

where Δ_{c}^{T} is the delta of the first OTM call and $\Delta_{p'}^{T}$ is equal to the delta of the first OTM put plus one (i.e. the delta of the corresponding in-the-money call). The variables $\sigma_{OTM,c}^{T}$ and $\sigma_{OTM,p}^{T}$ refer to the implied volatility of the first OTM call and of the first OTM put, respectively. The above methodology is then repeated using options expiring in the second-nearest maturity T', to extract $\sigma_{ATM}^{T'}$.

$$\sigma_{ATM}^{T'} = (\frac{\Delta_{p}^{T'} - 0.5e^{-qT}}{\Delta_{p}^{T'} - \Delta_{c}^{T'}})\sigma_{OTM,c}^{T'} + (\frac{0.5e^{-qT} - \Delta_{c}^{T'}}{\Delta_{p'}^{T'} - \Delta_{c}^{T'}})\sigma_{OTM,p}^{T'} \tag{4.14}$$

The ATM volatility $\sigma_{ATM}^{T_S}$ of a standardized 30-day synthetic option is obtained by linearly interpolating between the ATM volatilities of the nearest and second-nearest expirations T and T', respectively, using (4.15)

$$\sigma_{ATM}^{T_S} = (\frac{T'-T_S}{T'-T})\sigma_{ATM}^{T} + (\frac{T_S-T}{T'-T})\sigma_{ATM}^{T'} \tag{4.15}$$

where T_S is the time-to-expiration that corresponds to 30 calendar days (or equivalently 22 trading days). For notational convenience, throughout the remaining of the analysis the 30-day standardized B&S implied volatility and variance will be denoted by σ_{BS} and var_{BS}, respectively.

4.4.2 Corrado and Su Implied Variance

Although the Black and Scholes methodology has been widely applied throughout the related literature, it relies on certain assumptions, the most restrictive of which is arguably the lognormality of asset returns. The Corrado and Su (1996) framework relaxes the lognormality restriction by modifying the original B&S formula to account for non-zero skewness and excess kurtosis in the underlying's distribution using a Gram-Charlier series expansion of the standard normal density function. More specifically, they define a density function *g(z)* which accounts for non-normal skewness and kurtosis, described by the following equation, where *n(z)* represents the standard normal density function and μ_n is the standardized coefficient of the n^{th} moment of the asset's returns distribution.

$$g(z) = n(z)[1 + \frac{\mu_3}{3!}(z^3 - 3z) + \frac{\mu_4 - 3}{4!}(z^4 - 6z^2 + 3)] \qquad (4.16)$$

Within this framework, a call option's price is expressed as the sum of the theoretical B&S price and two correction terms related to non-normal skewness and excess kurtosis. Then, the Corrado and Su (C&S) call option price C_{CS} is given by:

$$C_{CS} = C_{BS} + \mu_3 M_3 + (\mu_4 - 3) M_4 \qquad (4.17)$$

with

$$M_3 = \frac{1}{3!} A_0 \sigma \sqrt{T}[(2\sigma\sqrt{T} - d_1)n(d_1) + \sigma^2 T N(d_1)]$$

$$M_4 = \frac{1}{4!} A_0 \sigma \sqrt{T}[(d_1^2 - 1 - 3\sigma\sqrt{T}(d_1 - \sigma\sqrt{T}))n(d_1) + \sigma^3 T^{\frac{3}{2}} N(d_1)]$$

where C_{BS} is the theoretical Black and Scholes call option price and M_3 and M_4 represent the marginal effects of non-normal skewness and kurtosis on the call's price, respectively. Moreover, theoretical C&S prices P_{CS} of put options are obtained by first estimating the price of the equivalent, same-strike call and then using the put-call parity:

$$P_{CS} = C_{CS} + Ke^{-rT} - A_0 \tag{4.18}$$

C&S implied volatilities σ_{CS}^{T} and $\sigma_{CS}^{T'}$, corresponding to the nearest and second-nearest expirations, respectively, are calculated by minimizing the following sums of squared errors between market prices Q_M and theoretical C&S prices Q_{CS} with respect to the moments vector (σ, μ_3, μ_4)

$$\min_{\sigma^T, \mu_3^T, \mu_4^T} SSE = \sum_{i=1}^{M} [(Q_{i,M}^{T} - Q_{i,CS}^{T})^2]$$

$$\min_{\sigma^{T'}, \mu_3^{T'}, \mu_4^{T'}} SSE = \sum_{i=1}^{M'} [(Q_{i,M}^{T'} - Q_{i,CS}^{T'})^2]$$

where M and M' are the number of observations for the maturities T and T', respectively. Standardized 30-day implied volatility σ_{CS} and implied variance var_{CS} are obtained by linearly interpolating across the nearest and second-nearest maturities. Finally, it should be noted that, in addition to assuming a more flexible distribution of asset returns, the Corrado and Su option pricing formula has the advantage of simultaneously estimating the next two higher implied moments, namely skewness and kurtosis, for which monthly estimates can be obtained through the same linear interpolation.

4.4.3 Model-Free Implied Variance

In recent years, the literature related to estimating implied variances has somewhat shifted focus from model-based estimation to model-free estimates. For instance, when calculating B&S and C&S implied volatilities, one jointly assumes the informational efficiency of the options market as well as the validity of the underlying model. Britten-Jones and Neuberger (2000) demonstrate that the future variance of asset returns can be calculated without the restriction of assuming a specific model for the returns distribution. In their influential paper, they demonstrate that volatility implied by options prices can be estimated as the expected sum of squared returns under the risk-neutral measure. More specifically, Britten-Jones and Neuberger show that the risk-neutral expected sum of squared returns in the time interval [0, T] is completely specified by a set of OTM options expiring at T.

$$E_0^Q[V_T] = E_o^Q[\int_0^T (\frac{dA_t}{A_t})^2] = 2e^{rT}[\int_0^{F_{0,T}} \frac{p(K,T)}{K^2} dK + \int_{F_{0,T}}^{\infty} \frac{c(K,T)}{K^2} dK] \tag{4.19}$$

where V_T is the integrated squared volatility of the asset, A_t is the asset's spot price at time t, and $F_{0,T}$ is the forward price at time 0 for delivery at time T. Moreover, $p(K,T)$ and $c(K,T)$ are the prices of OTM put and call options, respectively, with strike K and expiring at T.

In order to derive equation (4.19), only the stochastic process of the underlying is assumed to be continuous. The Britten-Jones and Neuberger method requires a continuum of strike prices K, with option prices quoted at every strike. In reality, however, empirical estimation of squared expected returns can only be done using a finite set of discrete strikes. Carr and Wu (2004) and Jiang and Tian (2005) relax the assumption of continuity and provide discrete versions of the Britten-Jones and Neuberger model. This study follows the methodology adopted, among others, by Taylor, Yadav and Zhang (2006) of using a finite set of OTM options written on an asset to estimate the asset's integrated variance until the options' expiration. The discrete version of (4.19) is then given as follows:[3]

$$\mathrm{var}_{MF} = \frac{2}{T} e^{rT} \sum_{i=1}^{M} \frac{\Delta K_i}{K_i^2} Q(K_i, T) - \frac{1}{T} [\frac{F_{0,T}}{K_0} - 1]^2 \tag{4.20}$$

where var_{MF} is the model-free expectation of the future variance, M is the number of strike prices used, and Q_i is the option's market price at strike K_i. Since K_0 denotes the strike price used to select either call or put options in the formula, the option price Q_i

[3] Taylor, Yadav and Zhang (2006) use the methodology adopted by the CBOE to calculate the model-free volatility expectation of the S&P 500 index over a standardized period of thirty days.

refers to calls when $K_i \geq K_0$, and to puts otherwise. Finally, ΔK_i is estimated as $\frac{K_{i+1} - K_{i-1}}{2}$.

As can be easily seen from the above equation, the value of $\frac{1}{T}[\frac{F_{0,T}}{K_0} - 1]^2$ depends on the strike K_0 that is chosen to select call or put prices in the summation term. However, this methodology uses a small number of actual option prices to infer a risk-neutral density and, therefore, to create a significantly large number of artificial option-strike combinations. This allows for K_0 to be set equal to $F_{0,T}$, so that the final term in equation (4.20) disappears. Consequently, in estimating var_{MF} only OTM calls and puts are used, with K_0 denoting the ATM strike price.

Although the discrete version of model-free implied variance in (4.20) can be empirically estimated, it is the case that a significantly large set of options is needed to obtain an accurate measure of the underlying's future variance. Since options are actually quoted at a relatively limited number of strikes, a fact that is especially pronounced for options on individual stocks, the methodology described by Malz (1997) is employed to construct implied volatility curves using a small set of market-traded options.

Within this framework, the implied volatility curve is fitted as a function of option deltas, as opposed to a function of option strikes. Malz (1997) argues that this methodology ensures that volatilities of options that are further from the money (OTM and ITM contracts) are grouped more closely together than those of near-the-money options. Taylor, Yadav and Zhang (2006) also mention that '... extrapolating a function of delta provides sensible limits for the magnitude of implied volatility curves'. Following this line of thought, a quadratic function of implied volatility is fitted with respect to option delta. In addition to capturing the 'volatility smile', the quadratic specification has the advantage of requiring a minimum of only three options to be estimated.

$$IV_i = \alpha + \beta\Delta_i + \gamma\Delta_i^2 \tag{4.21}$$

Equation (4.21) describes the quadratic function used to construct the implied volatility curve, where IV_i is the implied volatility of option i, and Δ_i is the option's delta.

Moreover, IV_i is the simple Black and Scholes implied volatility of option i, while Δ_i is the sensitivity of call option i to changes in the value of the underlying, measured in (4.11) and (4.12) as the first derivative of the Black and Scholes formula with respect to changes in the underlying's price A_0. It should be noted that, when calculating the model-free implied variance, option deltas are expressed as a function of σ^*, which is a constant measure of volatility used across all options (see also Bliss and Panigirtzoglou (2002) and Taylor, Yadav and Zhang (2006) for the use of σ^*), with the respective d_1 estimated as:

$$d_1(K_i) = \frac{\ln(\frac{A_0}{K_i}) + \frac{\sigma^* T^2}{2}}{\sigma^* \sqrt{T}}$$

Call deltas range from zero for deep OTM contracts with high strikes to e^{-rT} for deep ITM ones with low strikes. The respective put deltas range from $-e^{-rT}$ (deep ITM puts with high strikes) to zero (deep OTM puts with low strikes).

The parameter vector $\Phi = [\alpha, \beta, \gamma]$ of the quadratic function is estimated by minimizing the weighted sum of squared differences between observed volatilities IV and fitted volatilities $\hat{IV}(\Delta_j, \Phi)$ with respect to Φ, as given in (4.22):

$$\min_{\Phi} \sum_{j=1}^{M} w_j [IV_j - \hat{IV}_j(\Delta_j, \Phi]^2 \tag{4.22}$$

where M is the number of observed strikes and w_j is the weight of option j's delta. The weight w_j of option j is equal to $\Delta_j(1 - \Delta_j)$ and the minimization is subject to the constraint of fitted volatilities being strictly positive, $\hat{IV}(\Delta_j, \Phi) > 0$. The above weighting scheme ensures that deviations of fitted volatilities from observed levels are more heavily penalized for the nearest-the-money options, i.e. calls with deltas close to 0.50, compared to further-from-the-money contracts. Placing more weight in near-the-money options is compatible with the stylized fact that these options are more heavily traded, thus reducing the effect of possible outliers of illiquid ITM and OTM contracts.

Figure 4.7

Spot Rate's Implied Volatility

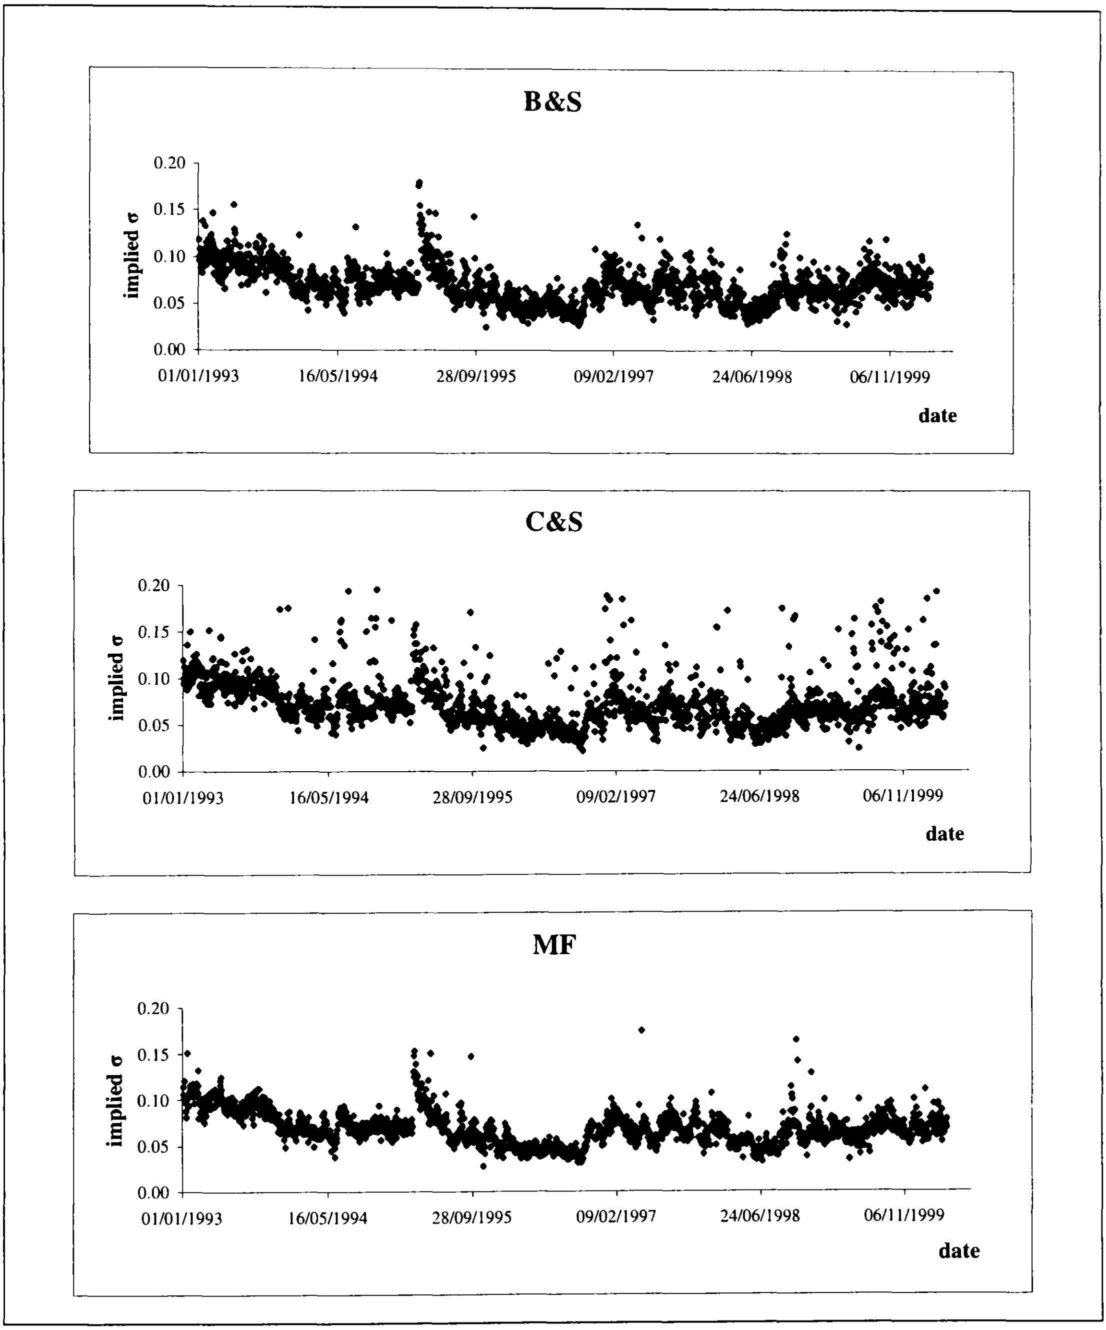

After fitting the implied volatility curve, a large set of artificial option prices is created using the vector Φ. More specifically, 1,000 equally spaced deltas ranging from 0 to e^{-rT} are used to extract the corresponding strikes. Then, option prices (both calls and puts) are estimated using the Black and Scholes formula with the respective combinations of strike price and volatility. Finally, the OTM contracts are identified and used in estimating the asset's integrated variance in equation (4.20), while standardized 30-day estimates are calculated by linearly interpolating between the nearest and second-nearest variances.[4]

4.5 OLS Estimation

4.5.1 Estimating the Standard Specification

Previous studies of the forward unbiasedness hypothesis have mainly focused on two types of specifications, based on exchange rate levels and exchange rate returns.[5] The first type is specified by a regression of future spot rate levels on current forward rates, described in (4.23), while the second specification involves regressing exchange rate returns on the forward premium and it is given in equation (4.24).

$$s_{t+\tau} = \alpha + \beta f_t^{t+\tau} + \varepsilon_t \tag{4.23}$$

$$s_{t+\tau} - s_t = \alpha + \beta(f_t^{t+\tau} - s_t) + \varepsilon_t \tag{4.24}$$

Similarly to the earlier stream of the literature on the forward premium puzzle, this study tests the standard specification where the forward rate is the only explanatory variable of the future spot rate. Before examining the effect of expected variance, the standard specification in (4.23) is estimated for the period 1993-2000 through OLS regressions, with Panel A of Table 4.2 presenting the results. Consistent with previous findings, the estimated slope for the entire sample is found to be less than the value of one that is predicted by UIP. However, at a level of 0.7690, the forward rate's beta is positive and statistically significant. Furthermore, the intercept is significantly positive (t-

---

[4] From the artificial set of option prices, the methodology uses OTM puts with strikes in the range $[0, F_{0,T})$ and OTM calls with strikes in the range $[F_{0,T}, \infty)$

[5] A third type refers to Error-Correction Models (ECM) which include lagged differences between the spot and the forward rates as additional explanatory variables of exchange rate returns.

stat = 18.77), indicating that changes in the forward rate fail to fully explain changes in future spot rates.

In addition to regressing equation (4.23) for the entire sample period, rolling 2-year estimations are performed in order to examine the evolution of observed betas across time. This step is motivated by the findings of Baillie and Bollerslev (2000), among other studies, which report a significant time-variation of the slope coefficient in forward premium regressions, indicating that the rejection of the forward unbiasedness hypothesis is probably dependant on the period examined rather than a universal characteristic of foreign exchange markets.

Table 4.2
OLS Results

Panel A: Standard Specification

$$s_{t+\tau} = \alpha + \beta f_t^{t+\tau} + \varepsilon_t$$

	α	β	Number of obs non-rejecting H_0: $\beta = 1$
	0.3666	0.7690	68
	(18.77)	(-18.77)	(4.30%)

Panel B: Extended JIT Specification

$$s_{t+\tau} = \alpha + \beta f_t^{t+\tau} + \gamma \operatorname{var}_t(s_{t+\tau}) + \varepsilon_t$$

	α	β	γ	Number of obs non-rejecting H_0: $\beta = 1$
var_{BS}	0.3935	0.7566	-1.3879	68
t-stat (= 0)	(19.65)	(-19.57)	(-5.36)	(4.30%)
t-stat (= -0.5)			(-3.43)	
var_{CS}	0.3676	0.7687	-0.0564	68
t-stat (= 0)	(18.78)	(-18.78)	(-0.82)	(4.30%)
t-stat (= -0.5)			(6.45)	
var_{MF}	0.3634	0.7717	-0.1255	66
t-stat (= 0)	(18.37)	(-18.11)	(-1.01)	(4.17%)
t-stat (= -0.5)			(3.01)	

This Table tabulates the results of OLS regressions of the forward rate unbiasedness hypothesis. Panel A refers to the standard specification while Panel B refers to the extended JIT specification.

As can be seem from Figure 4.8, rolling estimates of the slope coefficient indeed exhibit significant variability throughout the period examined. Starting from a level of around 0.60, beta follows an upward trend until roughly the middle of the sample period, when it experiences a sharp decline. The slope then slowly returns to 0.60 by 2000. Despite the significant volatility of the estimated slope, which ranges from a minimum of -0.06 to a maximum of 1.29, with the exception of a brief period in 1999 beta remains positive for the most part of the sample period. However, only 68 (4.30%) of the coefficients are statistically indistinguishable from the theoretical value of unity at the 5% confidence level, highlighting a significant deviation from the predictions of the UIP.

Figure 4.8

Rolling OLS Estimations under the Standard Specification

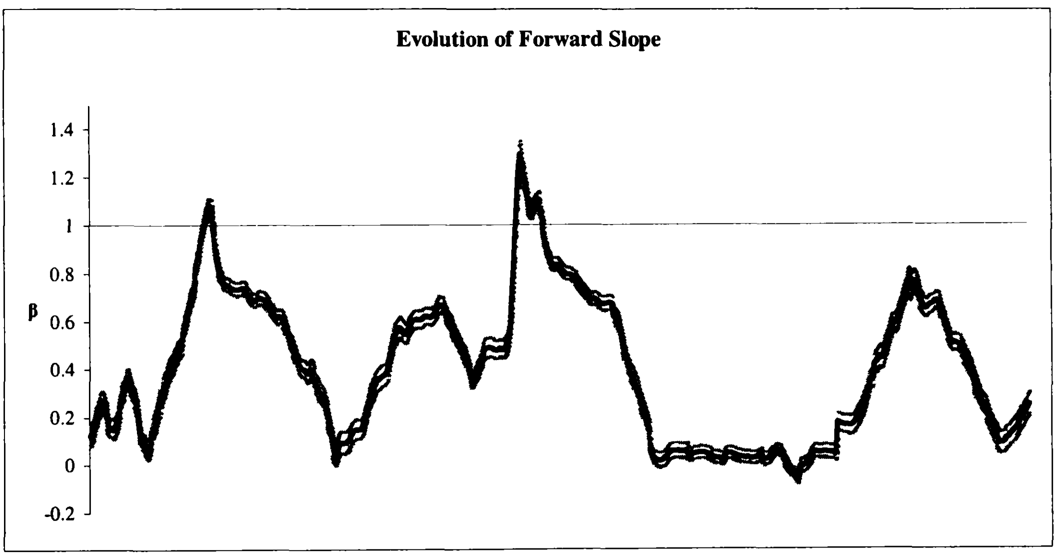

4.5.2 Accounting for the Future Variance of the Spot Rate

As has been discussed in Section 4.2, the specification in (4.23) fails to incorporate the JIT correction term for the expected variance of the future spot rate, considering the risk-premium to be equal only to the difference between $f_t^{t+\tau}$ and $s_{t+\tau}$. Although previous studies have shown that the contribution of the JIT term in explaining violations of UIP is not significant, it is the case that estimating $\mathrm{var}_t(s_{t+\tau})$, conditional on information available at time t, is not straightforward. Therefore, its documented inability to account for deviations of the forward rate's slope from its theoretical value of unity might be attributed, at least partly, to measurement error rather than to a fundamental quality.

This Section attempts to address the above concern by using option-implied variances as an alternative proxy for the exchange rate's future variance and to re-examine the forward unbiasedness hypothesis by testing the extended version in equation (4.25). Three different measures of option-implied variance are used, namely var_{BS}, var_{CS} and var_{MF}, corresponding to the three methodologies described in Section 4.4.

$$s_{t+\tau} = \alpha + \beta f_t^{t+\tau} + \gamma \operatorname{var}_t(s_{t+\tau}) + \varepsilon_t \qquad (4.25)$$

When estimating the OLS regressions, the three parameter vector [α, β, γ] is simultaneously estimated, instead of restricting γ to its theoretical value of -0.5, in order to allow for a more flexible framework. As can be seen from Panel B of Table 4.2, all three proxies for the future variance of the spot rate fail to improve the predictive power of the forward rate within the OLS framework. More specifically, when the extended specification in (4.25) is estimated for the entire sample period, the resulting forward betas are in fact slightly lower than the one estimated in (4.23) which, combined with the similar magnitude of standard errors, suggests an even further deviation from UIP. Moreover, although the estimates of the variance's slope are negative across all three proxies, γ is significant only for var_{BS}. Finally, with respect to the rolling 2-year estimations, introducing $var_t(s_{t+\tau})$ into the equation does not result in a higher proportion of individual estimates not rejecting the unbiasedness hypothesis as might have been expected. The first two proxies var_{BS} and var_{CS} produce the same proportion of betas lying within two standard errors of unity as that of the standard specification (4.30%), while the third proxy var_{MF} surprisingly results into a lower proportion of betas that are indistinguishable from one (4.17%).

Figure 4.9
Rolling OLS Estimations under the Extended Specification

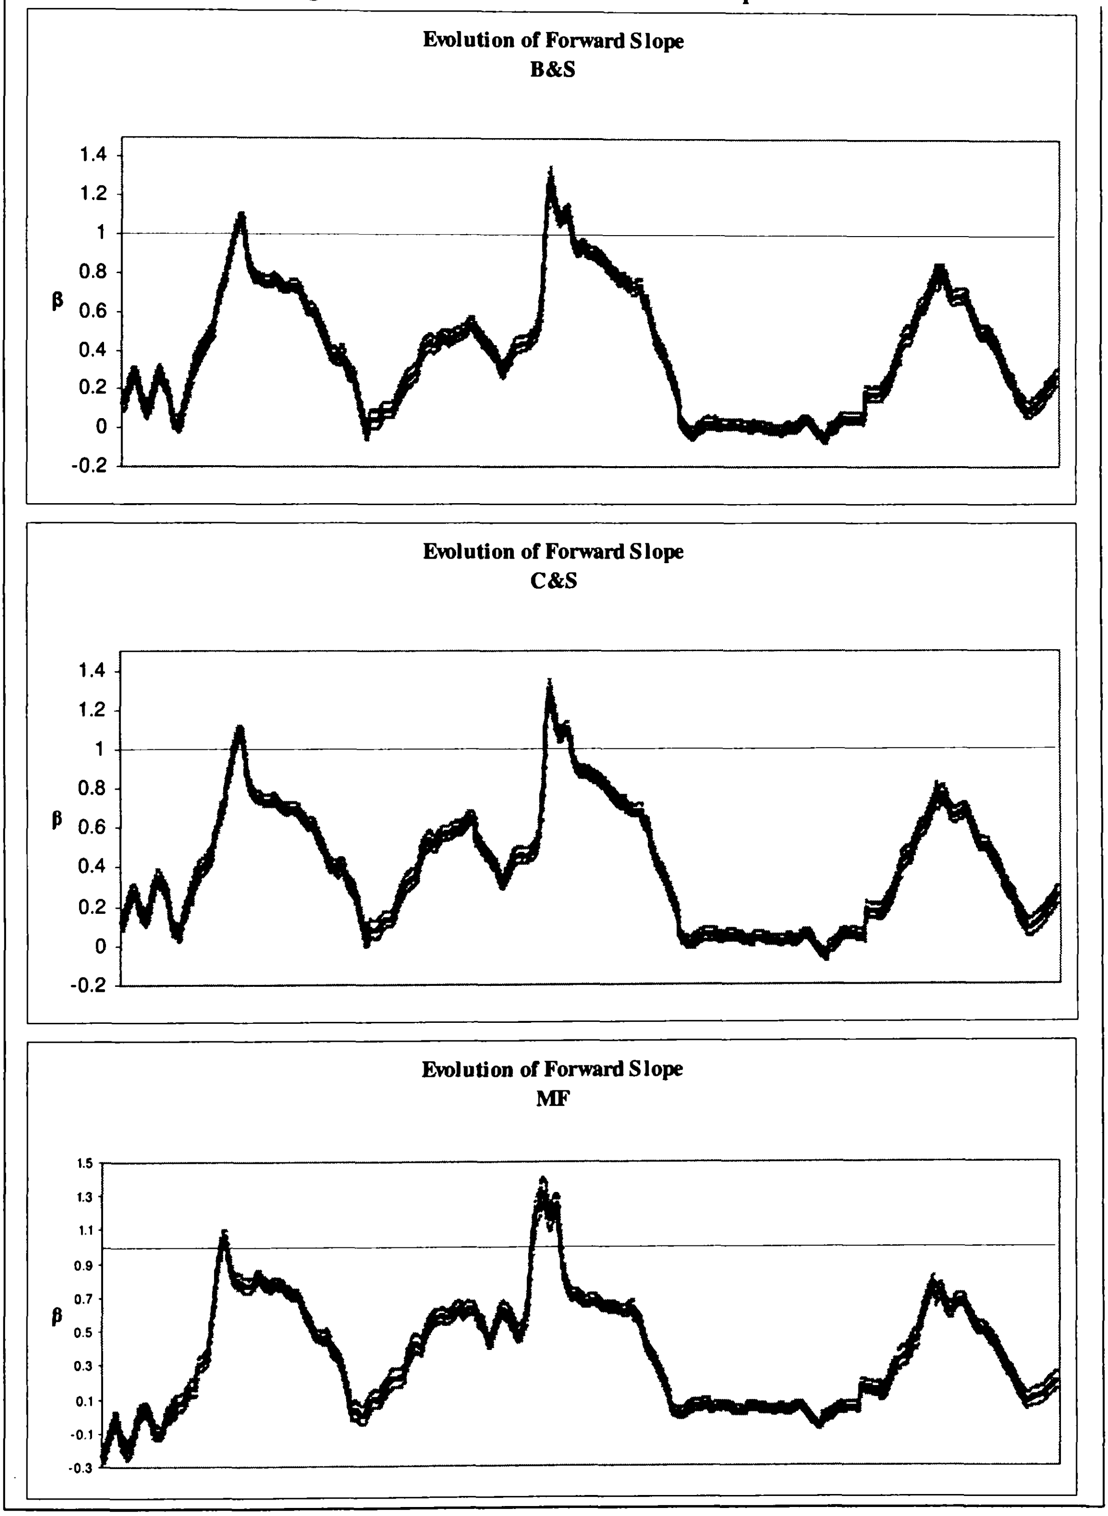

Overall, results from testing the forward unbiasedness hypothesis of the pound/dollar exchange rate using Ordinary Least Squares are in line with more recent empirical findings in the literature. Although forward betas are generally lower than their theoretical value of one, deviations from the UIP are less severe compared to earlier studies, as demonstrated by the fact that negative betas are obtained only for a small part of the period examined. In addition, incorporating the JIT term of the future spot rate's variance into the specification does not appear to have an impact on the predictive power of the forward rate over future spot rates.

4.6 FM-LAD Estimation

4.6.1 Complications from Using OLS Regressions

As has been mentioned in the previous Section, studies of the forward premium puzzle have mainly examined two types of specifications, involving exchange rate levels or exchange rate returns. It has been argued that neither type is free of certain limitations regarding the stationarity of the variables of interest, an issue that might be even more significant under the assumptions adopted when using the Ordinary Least Squares estimator to test the validity of UIP.

The first type is given by equation (4.23). A common concern with respect to this model refers to the order of cointegration between the dependant and the explanatory variables. More specifically, as has been discussed in Section 4.3, spot levels of the pound/dollar exchange rate were found to follow a unit root process while forward rates appear to be described by a fractionally integrated I(d) process, with d close to one.

Although fractional cointegration between spot and forward rates still produces consistent estimates of the forward's slope β, it is likely to complicate the statistical inference of its significance. For instance, Maynard and Phillips (2001) argue that the relatively close fit that the specification in (4.23) often produces may be due to the fact that it is '... simply reproducing the CIP relation, albeit with the interest differential hidden in the residual'. In addition, the finding that regressing future spot rates on current spot levels (see equation (4.26)) produces an even tighter fit compared to using forward rates is particularly puzzling.

$$s_{t+\tau} = \alpha + \beta s_t + \varepsilon_t \tag{4.26}$$

The second specification in (4.24) uses the forward premium/discount as the sole explanatory variable of exchange rate returns. However, similarly to the previous model, the orders of integration of $(s_{t+\tau} - s_t)$ and $(f_t^{t+\tau} - s_t)$ are not consistent, with spot returns found to be stationary and the forward premium rejecting the null of stationarity. Complications, therefore, arise when a short-memory variable is attempted to be explained by a long-memory regressor with a stochastic trend. Maynard and Phillips (2001) in particular argue that, unless the slope coefficient β is zero, the stochastic trend would be transferred to the spot return as well.

4.6.2 The FM-LAD Estimator

Given the above limitations regarding the OLS assumptions, the forward unbiasedness hypothesis is tested using the Fully Modified Least Absolute Deviations (FM-LAD) estimator derived in Phillips (1995), as an extension to the Least Absolute Deviations (LAD) technique. In addition to accommodating non-stationary data, endogeneity[6] and serially correlated errors, the FM-LAD methodology is also considered to be robust to outliers in the regressors and in the errors, which are documented characteristics of foreign exchange market models.[7] Finally, as has been argued by Phillips, McFarland and McMahon (1996), this technique allows for the use of exchange rate levels in (4.23) rather than their first differences in (4.24).

The FM-LAD estimator is a semi-parametric technique that treats the regression parameters α and β in (4.23) parametrically while treating nuisance parameters in a non-

[6] Endogeneity in forward unbiasedness regressions typically refers to interest rates and exchange market shocks being correlated. This monetary policy endogeneity has been shown to be more pronounced for short horizons, with the endogeneity bias of the forward unbiasedness regressor declining in longer horizons where macroeconomic fundamentals become more important (see Meredith and Ma (2002) for a more detailed analysis of the role of endogeneity in the forward premium puzzle). This phenomenon has been frequently referred to as 'leaning against the wind', where a central bank buying currency is associated with that currency depreciating.

[7] The FM-LAD estimator has been used in studies examining foreign exchange markets by Felmingham and Leong (2005) and Phillips, McFarland and McMahon (1996).

parametric way. The following discussion on the FM-LAD properties draws from the comprehensive analysis in Phillips, McFarland and McMahon (1996). First, consider the cointegrated system

$$y_t = \beta' x_t + u_{0t} \tag{4.27a}$$

$$\Delta x_t = u_{xt} \tag{4.27b}$$

where $u_t' = (u_{0t}, u_{xt}')$ is a stationary vector with length $m = (1 + m_x)$ and spectral density matrix $f_{uu}(\lambda)$. The long-run covariance matrix of u_t is given by:

$$\Omega_{uu} = 2\pi f_{uu}(0) = \begin{pmatrix} \Omega_{00} & \Omega_{0x} \\ \Omega_{x0} & \Omega_{xx} \end{pmatrix} \tag{4.28}$$

where the partition is conformable with that of u_t and $\Omega_{xx} > 0$, so that the number of unit roots in the stochastic process x_t is equal to the dimension of x_t (i.e. x_t is a full rank I(1) process). Also define the long-run covariance matrix of $w_t' = (\upsilon_t, u_{xt}')$ as

$$\Omega_{ww} = 2\pi f_{ww}(0) = \begin{pmatrix} \Omega_{uu} & \Omega_{ux} \\ \Omega_{xu} & \Omega_{xx} \end{pmatrix} \tag{4.29}$$

partitioned conformably with w_t. Note that υ_t (but not necessarily u_{xt}) is bounded and has finite moments of all orders. The LAD estimator of β in (4.27a) is then given by:

$$\beta_{LAD} = \arg\min_{\beta} \left\{ n^{-1} \sum_{1}^{n} \left| y_t - x_t'\beta \right| \right\} \tag{4.30}$$

where n denotes the sample size. The β_{LAD} estimator has an asymptotic normal distribution and is $\sqrt{n}$ consistent for β. However, when x_t follows a unit root and (4.27b) holds, the LAD estimator suffers from bias and non-scale nuisance parameter problems, similarly to OLS. In order to account for these limitations, Phillips (1995) designs the FM-LAD estimator by modifying LAD to account for endogeneity in the regressor variables and for serial dependence in the errors. The FM-LAD estimator is defined by:

$$\beta_{LAD}^{+} = \beta_{LAD} - (\frac{1}{2}\hat{f}(0))(X'X)^{-1}(X'\Delta X - n\hat{\Delta}_{xv}^{+}) \tag{4.31}$$

where $X'X = \sum_{1}^{n} x_t x_t'$, $X'\Delta X = \sum_{1}^{n} x_t \Delta x_t'$ and $\hat{f}(0)$ is a consistent estimator of the probability density of u_{0t} at the origin. Also, $\hat{\Delta}_{xv}^{+}$ is a consistent estimator of the one-sided long-run covariance matrix

$$\Delta_{xv}^{+} = \sum_{k=0}^{\infty} E(u_{x0} v_k^{+}) \tag{4.32}$$

where

$$v_t^{+} = v_t - \Omega_{vx}\Omega_{xx}^{-1}\Delta x_t \tag{4.33}$$

and

$$v_t = sign(u_{0t}) \tag{4.34}$$

The error v_t and the modified error v_t^{+}, are estimated through a first-stage LAD regression which produces the error estimate $\hat{u}_{0t} = y_t - \beta_{LAD}' x_t$ and $\hat{u}_t = sign(\hat{u}_{0t})$. Then Δ_{xv}^{+} can be estimated and, using conventional kernel estimates of the long-run covariance matrices Ω_{vx} and Ω_{xx}, $\hat{v}_t^{+}$ is constructed, with

$$\hat{v}_t^{+} = \hat{v}_t - \hat{\Omega}_{0x}\hat{\Omega}_{xx}^{-1}\Delta x_t \tag{4.35}$$

Moreover, Δ_{xv}^{+} can be rewritten as

$$\Delta_{xv}^{+} = \Delta_{xv} - \Delta_{xx}\Omega_{xx}^{-1}\Omega_{xv} \tag{4.36}$$

where $\Delta_{xv} = \sum_{k=0}^{\infty} E(u_{x0}v_k)$, $\Delta_{xx} = \sum_{k=0}^{\infty} E(u_{x0}u'_{xk})$

so that the estimation of Δ_{xv}^{+} involves the estimation of the four submatrices $\Delta_{xv}, \Delta_{xx}, \Omega_{xx}$, and Ω_{xv}. Phillips (1995) shows that when system (4.27) has finite variance errors, the β_{LAD}^{+} estimator in (4.31) is asymptotically mixed normal, i.e.

$$(\beta_{LAD}^{+} - \beta) \sim N(0, (\frac{1}{2}f(0))^2[\omega_{vv\cdot x} \otimes (X'X)^{-1}]) \tag{4.37}$$

where $\omega_{vv\cdot x} = \Omega_{vv} - \Omega_{vx}\Omega_{xx}^{-1}\Omega_{xv}$. Also note that the asymptotic mixed normal approximation in (4.37) holds even if the error variances in system (4.27) are not finite, making the FM-LAD estimator a flexible testing technique in the case of foreign exchange markets.

4.6.3 FM-LAD Results

4.6.3.1 The Standard Specification

Panel A of Table 4.3 presents the results from the FM-LAD estimation of the forward unbiasedness hypothesis in (4.23).[8] When future spot rates are regressed on current forward rates for the entire sample period, a slope coefficient of 0.8342 is obtained. This value of β is significantly positive and closer to one compared to the OLS estimate, albeit statistically different from the value of unity that UIP predicts. The above result is in line with findings from more recent studies which suggest that deviations from UIP have been relatively less severe in the 90's compared to previous periods. However, the statistically significant intercept term seems to indicate that forward rates fail to adequately predict subsequent realizations of spot rates.

[8] The FM-LAD estimation was programmed using the R-software.

Moreover, β starts highly positive at a level of around 0.60 at the beginning of the period and fluctuates between 0.60 and 0.90 until late in 1997. However, β experiences a large drop during the later part of the sample period, reaching its lowest level of 0.13 midway through 1998, and slowly approaching 0.50 in the early 2000. Finally, considering the slope's standard error, an easy way of summarizing the predictive power of $f_t^{t+\tau}$ over $s_{t+\tau}$ across time is by examining the number of individual rolling estimations which produce betas that lie less than two standard errors of their theoretical value of unity. Regressing the standard model in (4.23), thus, results in 55 estimated betas (4.12% of the overall set of coefficients) being statistically insignificant from one at the 5% confidence interval, with these estimates generally found during 1995.

Figure 4.10
Rolling FM-LAD Estimations under the Standard Specification

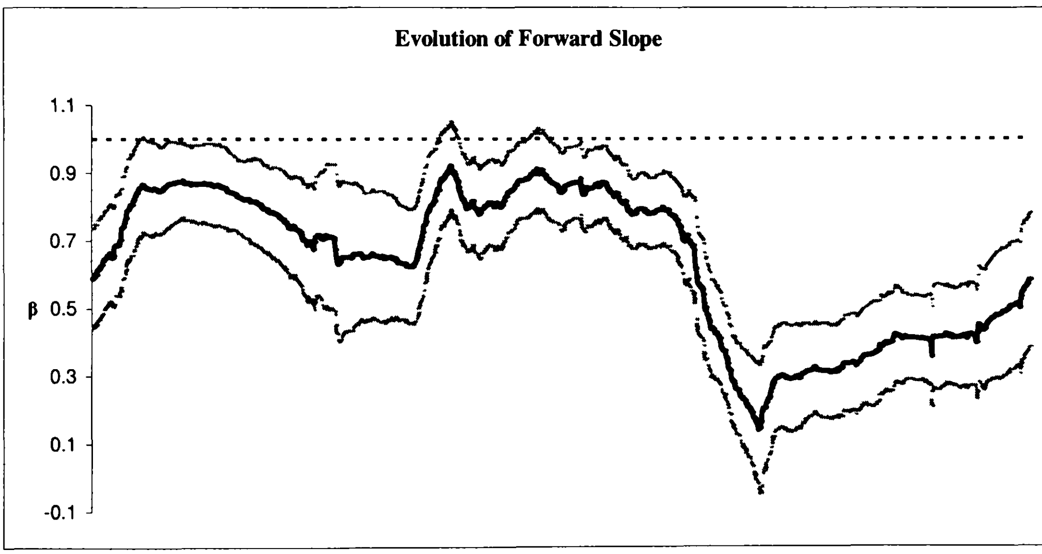

4.6.3.2 The Extended Specification

Considering the point estimates for the entire sample period, the inclusion of $\mathrm{var}_t(s_{t+\tau})$ as a correction term appears to have little contribution (if any) in obtaining a β coefficient closer to unity, irrespective of the proxy used. More specifically, including var_{CS} and var_{MF} as additional regressors results in only a marginal increase in the forward's slope from 0.8342 to 0.8363 and 0.8433, respectively, while the var_{BS} term actually produces a slightly lower β (0.8315) compared to the standard specification. Furthermore, the coefficient γ of the expected spot variance is significant only in the case of the B&S estimate (t-stat = -2.26), marginally significant for var_{CS} (t-stat = -1.67), and insignificant for var_{MF} (t-stat = -1.20). Finally, the intercepts remain significant and at levels comparable to those previously reported for (4.23).

Table 4.3
FM-LAD Estimator Results

Panel A: Standard Specification

$$s_{t+\tau} = \alpha + \beta f_t^{t+\tau} + \varepsilon_t$$

	α	β	Number of obs non-rejecting H_0: $\beta = 1$
	0.0766	0.8342	55
	(5.76)	(-5.76)	(4.12%)

Panel B: Extended JIT Specification

$$s_{t+\tau} = \alpha + \beta f_t^{t+\tau} + \gamma \operatorname{var}_t(s_{t+\tau}) + \varepsilon_t$$

	α	β	γ	Number of obs non-rejecting H_0: $\beta = 1$
var_{BS}	0.0875	0.8315	-0.1360	160
t-stat (= 0)	(5.97)	(-5.85)	(-2.26)	(11.99%)
t-stat (= -0.5)			(6.05)	
var_{CS}	0.0795	0.8363	-0.0503	139
t-stat (= 0)	(5.90)	(-5.77)	(-1.67)	(10.42%)
t-stat (= -0.5)			(14.93)	
var_{MF}	0.0767	0.8433	-0.0462	167
t-stat (= 0)	(5.70)	(-5.29)	(-1.20)	(12.52%)
t-stat (= -0.5)			(11.79)	

This Table tabulates the results of FM-LAD regressions of the forward rate unbiasedness hypothesis. Panel A refers to the standard specification while Panel B refers to the extended JIT specification.

Figure 4.11
Rolling FM-LAD Estimations under the Extended Specification

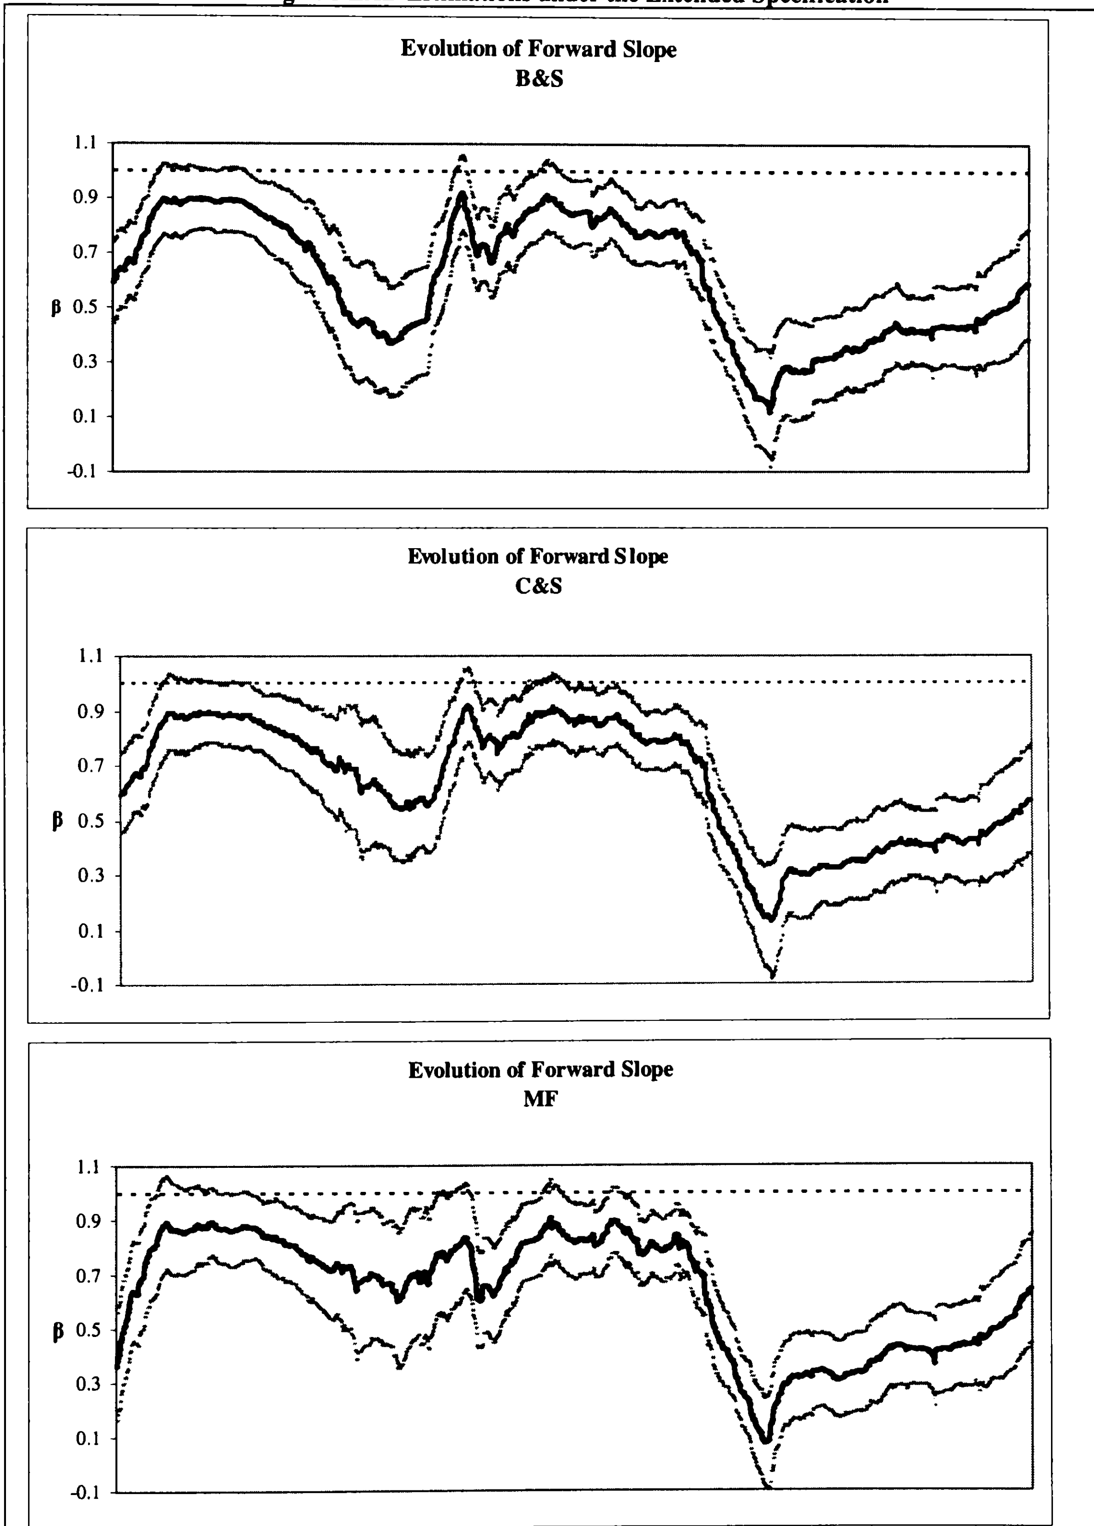

However, when rolling 2-year regressions are considered, adding the JIT term in the specification appears to have a significant impact on the forward rate's explanatory power of future spot rates. As can be seen from Figure 4.11, β estimates exhibit variability through time, while following a similar pattern to that of the standard specification in Figure 4.10. More importantly, though, the 5% confidence bands have moved closer to unity in the case of all three $\mathrm{var}_t(s_{t+\tau})$ proxies, resulting in more coefficients failing to reject the unbiasedness hypothesis of $\beta = 1$. The proportion of observations with forward slopes that are statistically indistinguishable from one is 11.99%, 10.42% and 12.52% for var_{BS}, var_{CS} and var_{MF}, respectively. This represents an increase by a factor of three compared the standard specification in (4.23) and goes some way into supporting the role of the JIT future spot variance in producing results that are more in line with the UIP predictions.

Overall, the results suggest that including an option-implied proxy for the future variance of the spot rate in an extended forward unbiasedness specification and using the FM-LAD estimator instead of the standard OLS framework is associated with a significantly higher proportion of betas that are in line with interest parity. It can be easily seen from comparing Figures 4.8, 4.10 and 4.11 that this improvement stems from two factors. First, the beta coefficients of the extended specification in (4.25) estimated through FM-LAD are systematically higher and, therefore, closer to the theoretical value of unity compared to betas obtained by estimating the standard specification in (4.23) through OLS. Also, the resulting standard errors are higher when the FM-LAD estimator is used on (4.25) compared to the OLS standard errors, making it more difficult to reject Uncovered Interest Parity. More specifically, when the observations are examined where parity is rejected under OLS/(4.23) but not rejected under FM-LAD/(4.25), all but three cases are found to be characterized by higher betas as well as by higher standard errors. However, these two effects do not have the same magnitude since the increase in the forward betas is typically higher than the respective increase in the regressions' standard errors. For instance, when the extended specification (4.25) is estimated under FM-LAD, standard errors increase on average by a factor higher than three (353%) relative to the average standard error of the standard specification (4.23) under OLS. Although the forward betas increase by less than that in percentages, their magnitude is obviously higher, making their absolute effect more significant in reducing the proportion of parity violations.

4.7 Conclusion

This Chapter has examined the forward premium anomaly, i.e. the widely reported finding that when spot rates are regressed on forward rates, the resulting slope coefficients deviate from one and, in many cases, fall below zero. A negative forward slope represents a significant violation of Uncovered Interest Parity and implies that, not only does the forward rate fail to predict the future level of exchange rates, but that it effectively predicts spot changes of the wrong sign. For instance, if the forward rate is higher than the current spot rate, a negative beta is associated with a future spot rate that, instead of appreciating to reach the forward level, it will actually tend to systematically depreciate and fall below its current level.

The empirical results confirm previous findings that using shorter samples in rolling regressions produces slope coefficients that are relatively dispersed, positive and, for some 2-year periods, statistically indistinguishable from unity. Furthermore, the FM-LAD estimator results in slopes that are more consistent with UIP compared to Ordinary Least Squares Estimates, since the former technique is able to address some common econometric concerns related to foreign exchange markets, such as fractional cointegration between spot and forward rates, serially correlated errors and endogeneity.

Finally, the future variance of spot returns appears to have some explanatory power over previously reported deviations from UIP. More specifically, adding the option-implied expected variance as an additional regressor in the forward unbiasedness specification results in a significantly higher proportion of slopes that are equal to their theoretical value of unity, providing support for the hypothesis that Jensen's Inequality is related to the magnitude of the risk-premium observed in foreign exchange.

Appendix A

Derivation of the Extended JIT Forward Unbiasedness Specification

A function $f(x)$ is convex if the expected value of the function is higher than the function of the expectation, as given in (A.1). An example of a convex function is $f(x) = \frac{1}{x}$.

$$E(f(x)) > f(E(x)) \tag{A.1}$$

On the other hand, a function $f(x)$ is concave if the expected value of the function is lower than the function of the expectation, as given in (A.2). An example of a concave function is the logarithmic function $f(x) = \ln(x)$.

$$E(f(x)) < f(E(x)) \tag{A.2}$$

For the logarithmic function, in particular, it can be shown that the logarithm of the expectation is higher than the expectation of the logarithm by half the log's variance:

$$\ln E(x) = E(\ln(x)) + \frac{1}{2}\operatorname{var}(\ln(x)) \tag{A.3}$$

Jensen's Inequality in foreign exchange markets stems from the fact that exchange rates are quoted as ratios of currencies and, given the convexity of $f(x) = \frac{1}{x}$ and equation (A.1), the expected rate of appreciation of one currency can never be equal to the expected rate of depreciation of the other, a fact that is also referred to as *Siegel's Paradox.* For instance, if the forward rate is assumed to be a conditionally unbiased predictor of the future spot rate, i.e. $F_t^{t+\tau} = E_t[S_{t+\tau}]$, *Siegel's Paradox* is expressed as follows:

$$E_t(\frac{1}{S_{t+\tau}}) > \frac{1}{E_t[S_{t+\tau}]} = \frac{1}{F_t^{t+\tau}} \tag{A.4}$$

Define $F_t^{RN,t+\tau}$ as the risk-neutral forward rate. The risk-neutral forward rate is given by:

$$F_t^{RN,t+\tau} = \frac{E_t[\frac{S_{t+\tau}}{P_{t+\tau}}]}{E_t[\frac{1}{P_{t+\tau}}]} \tag{A.5}$$

Furthermore, assume rational expectations and that risky assets are priced with a stochastic discount factor (pricing kernel). Under risk-neutrality there are no expected real profits from forward market speculation, i.e. the risk-neutral investor would arbitrage the market until the following condition is met (Engel (1984) provides a more detailed discussion of (A.6) in a single consumption good economy, and Engel (1992) shows that this equation still describes the risk-neutral forward rate when utility is a function of more than one good):

$$E_t[\frac{F_t^{RN,t+\tau} - S_{t+\tau}}{P_{t+\tau}}] = 0 \tag{A.6}$$

Using the concavity property of the logarithmic function in (A.3) and assuming that exchange rates and prices levels are log-normally distributed, equation (A.6) can be rewritten as:

$$\ln(E_t[\frac{F_t^{RN,t+\tau} - S_{t+\tau}}{P_{t+\tau}}]) = E_t[\ln(\frac{F_t^{RN,t+\tau} - S_{t+\tau}}{P_{t+\tau}})] + \frac{1}{2}\mathrm{var}(\ln(\frac{F_t^{RN,t+\tau} - S_{t+\tau}}{P_{t+\tau}})) = 0 \tag{A.7}$$

and, finally, as:

$$E_t[s_{t+\tau}] - f_t^{RN,t+\tau} = -\frac{1}{2}\text{var}_t(s_{t+\tau}) + \text{cov}_t(s_{t+\tau}, p_{t+\tau}) \tag{A.8}$$

where the lower case variables $s_{t+\tau}$, $f_t^{RN,t+\tau}$ and $p_{t+\tau}$ correspond to the logarithmic values of the upper case variables $S_{t+\tau}$, $F_t^{RN,t+\tau}$ and $P_{t+\tau}$, respectively (see also equation (4.4) in the main text).

Moving from risk-neutral to real-world expectations, define $F_t^{RW,t+\tau}$ as the real-world forward rate. Then, if utility is time-separable with a constant rate of time-preference, the equivalent of (A.6) is given by:

$$E_t[\frac{F_t^{RW,t+\tau} - S_{t+\tau}}{P_{t+\tau}} \times \frac{\beta u'(C_{t+\tau})}{u'(C_t)}] = 0 \tag{A.9}$$

where β is the stochastic discount factor, $u(\cdot)$s the utility function and the prime (') indicates the first derivative. Also, $A_{t+\tau}$refers to the intertemporal marginal rate of substitution:

$$A_{t+\tau} \equiv \frac{\beta u'(C_{t+\tau})}{u'(C_t)} \tag{A.10}$$

Similarly to the above, assuming all variables are log-normally distributed and defining $\alpha_{t+\tau}$ as the logarithm of $A_{t+\tau}$, equation (A.9) can be rewritten as:

$$E_t[s_{t+\tau}] - f_t^{RW,t+\tau} = -\frac{1}{2}\text{var}_t(s_{t+\tau}) + \text{cov}_t(s_{t+\tau}, p_{t+\tau}) - \text{cov}_t(s_{t+\tau}, \alpha_{t+\tau}) \tag{A.11}$$

References

Azar, Samih Antoine, 2008, Jensen's Inequality in Finance, *International Atlantic Economic Society* 14, 433-440

Bacchetta, Philippe, and Eric van Wincoop, 2007, Random Walk Expectations and the Forward Discount Puzzle, *American Economic Review, Papers and Proceedings* 97, 346-350

Bacchetta, Philippe, and Eric van Wincoop, 2008, Incomplete Information Processing: A Solution to the Forward Discount Puzzle, *Study Center Gerzensee working paper*

Backus, David, Allan Gregory, and Chris Telmer, 1993, Accounting for Forward Rates in Markets for Foreign Currency, *Journal of Finance* 48, 1887-1908

Bahmani-Oskooee, Mohsen, and Satya Das, 1985, Transaction Costs and the Interest Parity Theorem, *Journal of Political Economy* 93, 793-799

Baillie, Richard, and Tim Bollerslev, 1990, A Multivariate Generalized GARCH Approach to Modeling Risk Premia in Forward Foreign Exchange Markets, *Journal of International Money and Finance* 9, 309-324

Baillie, Richard, and Tim Bollerslev, 1994, The Long Memory of the Forward Premium, *Journal of International Money and Finance* 13, 565-571

Bekaert, Geert, and Robert Hodrick, 1993, On Biases in the Measurement of Foreign Exchange Risk Premiums, *Journal of International Money and Finance* 12, 115-138

Bernoth, Kerstin, Juergen von Hagen, and Casper de Vries, 2007, The Forward Premium Puzzle: New Evidence from Futures Contracts, *DNB working paper*

Black, Fisher, and Myron Scholes, 1973, The pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637-654

Bliss, Robert, and Nikolaos Panigirtzoglou, 2002, Testing the Stability of Implied Probability Density Functions, *Journal of Banking and Finance* 26, 381-422

Boyer, Russell, and Charles Adams, 1988, Forward Premia and Risk Premia in a Simple Model of Exchange Rate Determination, *Journal of Money, Credit and Banking* 20, 633-644

Boyrie, Maria, Yong Kim, and Simon Pak, 2006, Price Risk and Bid-Ask Spreads of Currency Options, *Derivatives Use, Trading and Regulation* 12, 115-125

Britten-Jones, Mark, and Anthony Neuberger, 2000, Option Prices, Implied Price Processes, and Stochastic Volatility, *Journal of Finance* 55, 839-866

Byers, David, and David Peel, 1991, Some Evidence on the Efficiency of the Sterling-Dollar and Sterling-Franc Forward Exchange Rates in the Interwar Period, *Economics Letters* 35, 317-322

Carr, Peter, and Liuren Wu, 2004, Time-Changes Levy Processes and Option Pricing, *Journal of Financial Economics* 71, 113-141

Chaboud, Alain, and Jonathan Wright, 2005, Uncovered Interest Parity, It Works, but not for Long, *Journal of International Economics* 66, 349-362

Cheng, Ai-Ru, Ronald Gallant, Chuanshu Ji, and Beom Lee, 2009, A Gaussian Approximation Scheme for Computation of Option Prices in Stochastic Volatility Models, forthcoming in the *Journal of Econometrics*

Chesney, Marc, and M. Jeanblanc, 2004, Pricing American Currency Options in an Exponential Levy Model, *Applied Mathematical Finance* 11, 207-225

Chinn, Menzie, 2006, The (Partial) Rehabilitation of the Interest Parity in the Floating Rate Era: Longer Horizons, Alternative Expectations, and Emerging Markets, *Journal of International Money and Finance* 25, 7-21

Chinn, Menzie, and Guy Meredith, 2004, Monetary Policy and Long Horizon Uncovered Interest Parity, *IMF Staff Papers* 51, 409-430

Clinton, Kevin, 1988, Transactions Costs and Covered Interest Arbitrage: Theory and Evidence, *Journal of Political Economy* 96, 358-370

Corrado, Charles, and Tie Su, 1996, Skewness and Kurtosis in S&P500 Index Returns Implied by Option Prices, *Journal of Financial Research* 19, 175-192

Crowder, William, 1994, Foreign Exchange Market Efficiency and Common Stochastic Trends, *Journal of International Money and Finance* 13, 551-564

Cumby, Robert, 1988, Is it Risk? Explaining Deviations from Uncovered Interest Parity, *Journal of Monetary Economics* 22, 279-299

Engel, Charles, 1984, Testing for the Absence of Expected Real Profits from Forward Market Speculation, *Journal of International Economics* 17, 299-308

Engel, Charles, 1992, On the Foreign Exchange Risk Premium in a General Equilibrium Model, *Journal of International Economics* 32, 305-319

Engel, Charles, 1996, The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence, *Journal of Empirical Finance* 3, 123-192

Engel, Charles, 1999, On the Forward Exchange Risk Premium in Sticky-Price General Equilibrium Models, *International Tax and Public Finance* 6, 491-505

Evans, Martin, 1995, Peso Problems: Their Theoretical and Empirical Implications, *Handbook of Statistics: Statistical Methods in Finance* (North-Holland, Amsterdam)

Felmingham, Bruce, and Susan Leong, 2005, Parity Conditions and the Efficiency of the Australian 90- and 180-day Forward Markets, *Review of Financial Economics* 14, 127-145

Froot, Kenneth, and Jeffrey Frankel, 1989, Forward Discount Bias: Is it an Exchange Risk Premium, *Quarterly Journal of Economics* 104, 139-161

Garman, Mark, and Steven Kohlhagen, 1983, Foreign Currency Option Values, *Journal of International Money and Finance* 2, 231-237

Geweke, John, and Susan Porter-Hudak, 1983, The Estimation and Application of Long Memory Time Series Models, *Journal of Time Series Analysis* 4, 221-238

Gourinchas, Pierre-Olivier, and Aaron Tornell, 2004, Exchange Rate Puzzles and Distorted Beliefs, *Journal of International Economics* 63, 303-333

Hodrick, Robert, 1987, The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets, (Harwood)

Hodrick, Robert, 1989b, U.S. International Capital Flows: Perspectives from Rational Maximizing Models, *Carnegie-Rochester Conference Series on Public Policy* 30, 231-288

Jiang, George, and Yisong Tian, 2005, Model-Free Implied Volatility and its Information Content, *The Review of Financial Studies* 18, 1305-1342

Lewis, Karen, 1994, Puzzles in International Financial Markets, *The Handbook of International Economics*, (North-Holland, Amsterdam)

Malz, Allan, 1997, Option-Implied Probability Distributions and Currency Excess Returns, *Federal Reserve Bank of New York Staff Report* 32, November 1997

Mark, Nelson, Yangru Wu, and Weike Hai, 1997, Understanding Spot and Forward Exchange Rate Regressions, *Journal of Applied Econometrics* 12, 715-734

Maynard, Alex, and Peter Phillips, 2001, Rethinking an Old empirical Puzzle: Econometric Evidence on the Forward Discount Anomaly, *Journal of Applied Econometrics* 16, 671-708

McCallum, Bennett, 1994, A Reconsideration of the Uncovered Interest Parity Relationship, *Journal of Monetary Economics* 33, 105-132

Meredith, Guy, and Yue Ma, 2002, The Forward Premium Puzzle Revisited, *IMF working paper* 1-39

Phillips, Peter, 1995, Robust Nonstationary Regression, *Econometric Theory* 11, 912-951

Phillips, Peter, James McFarland, and Patrick McMahon, 1996, Robust Tests of Forward Exchange Market Efficiency with Empirical Evidence from the 1920s, *Journal of Applied Econometrics* 11, 1-22

Pong, Shiuyan, Mark Shackleton, Stephen Taylor, and Xinzhong Xu, 2004, Forecasting Currency Volatility: A Comparison of Implied Volatilities and AR(FI)MA Models, *Journal of Banking and Finance* 28, 2541-2563

Soderlind, Paul, and Lars Svensson, 1997, New Techniques to Extract Market Expectations from Financial Instruments, *Journal of Monetary Economics* 40, 383-429

Taylor, Stephen, Pradeep Yadav, and Yuanyuan Zhang, 2006, The Information Content of Implied Volatilities and Model-Free Volatility Expectations: Evidence from Options Written on Individual Stocks, *working paper*, Lancaster University

Chapter 5
An Examination of the Efficiency of Emerging Options Markets: The Case of the Athens Derivatives Exchange

5.1 Introduction

5.1.1 Literature Review

Since Sharpe (1964) and Lintner (1965) introduced the Capital Asset Pricing Model (CAPM), there has been a wide body of literature examining its empirical performance in pricing various classes of financial assets. Within the CAPM's mean-variance framework, investors are only interested in a security's contribution to their portfolio's systematic variance. In other words, assets with returns that exhibit a higher covariance with market returns will add to the overall portfolio's riskiness, therefore requiring higher expected returns compared to their lower covariance counterparts. Despite its subsequent popularity, though, many studies have challenged the validity of the CAPM, concluding that its single pricing factor (namely the asset's covariance with the market as measured by its beta) might be insufficient in explaining expected returns.

Fama and French (1992), motivated by the systematic variance's apparent lack of explanatory power over expected returns, attempt to improve the CAPM by incorporating two additional factors. The three Fama and French (FF) factors include the original model's market risk premium, the returns of a portfolio of small minus large capitalization stocks, and those of a portfolio of value minus growth stocks. The FF factors are based on the empirical observation that some classes of stocks systematically outperform the market, even after controlling for these assets' betas, and, despite the lack of a theoretical framework supporting the inclusion of these factors, FF find that their extended model performs much better than the traditional CAPM in explaining the cross-section of equity returns.

One explanation that has been proposed for the above empirical finding is the possibility that the FF factors are acting as proxies for the asset's higher systematic moments. More specifically, whereas in the CAPM framework investors form portfolios by examining only the incremental effect of adding another security in the portfolio's total systematic variance, in reality they are likely to be also interested in higher systematic moments of the asset's distribution. Since the non-normality of asset returns is a well documented empirical finding, asset pricing models that are based on the assumption of normally distributed returns are likely to ignore important risk factors that are being priced by the market.

Following this line of thought, some studies have examined whether the next two systematic moments are priced in addition to systematic variance. Kraus and Litzenberger (1976) provide one of the first empirical papers on risk-factors associated with higher co-moments by testing an extended version of the CAPM which includes an additional factor to account for the coskewness of asset returns with the market. Their results indicate that the new model is better in explaining the cross-section of equity returns, since systematic skewness appears to be priced, with the market assigning a negative risk-premium to positive levels of the third systematic moment. Furthermore, Fang and Lai (1997) introduce the fourth systematic moment into the analysis. Testing a three-moment model, they find that cokurtosis is priced in their cross-section of equity returns in addition to covariance and coskewness (see also Christie-David and Chaudhry (2001) for a similar examination of the futures market).

Despite the large number of past studies that examine the performance of pricing models in the case of equity, there has been relatively limited interest in expected option returns. Since options are risky assets, standard capital asset pricing theory predicts that they should earn a risk-premium related to the systematic risk they are exposed to. Coval and Shumway (2001) further demonstrate that, under a set of realistic assumptions, option returns must be increasing in strike price space, while calls should earn a return in excess of that of the underlying asset and puts should have an expected return below the risk-free rate. Focusing on calls and puts written on the S&P 500 index between 1990 and 1995, they find that option returns in their sample indeed exhibit the above mentioned characteristics. However, returns do not appear to vary linearly with their respective market betas, indicating that additional factors are potentially priced.

In contrast, Ni (2006) finds that the Coval and Shumway (2001) theoretical predictions do not apply for call options written on individual stocks. Examining a sample of US calls for the period 1996-2005, she reports average call option returns that are decreasing in strike price, with OTM calls earning negative returns. This puzzling finding is potentially explained through investors' seeking of idiosyncratic skewness, leading to higher than expected OTM call prices and, hence, lower returns. Jones (2001) analyzes a set of S&P 500 index options and concludes that idiosyncratic variance alone is insufficient in explaining short-term OTM put returns. He argues, therefore, that a multi-factor model is necessary to understand the risk-premia associated with options. Moreover, Broadie, Chernov and Johannes (2009) examine a larger sample of S&P 500 index options, namely from 1987 to 2005, with particular emphasis on puts. Contrary to Coval and Shumway, they report that the Black and Scholes (1976) option pricing model cannot be rejected based on deep OTM put returns. Also, ATM put and straddle returns are found to be consistent with jump models.

Liu (2007) focuses on arguably the two most common sources of risk in the options market, namely changes in the value of the underlying and changes in the underlying's volatility. By forming delta and vega neutral straddles with options written on the FTSE100, she explores the hypothesis that options portfolios that are immune to delta and vega risk should earn the risk-free rate, and finds that this prediction is supported for ATM and ITM portfolios. However, OTM straddles appear to earn significantly negative returns, with one potential explanation for this result being the fact that delta and vega neutrality, measured as Black and Scholes local sensitivities, do not necessarily hold for the entire holding period of the straddles. The paper also examines risk-reversals, which are option positions that profit from negative skewness, and finds that, even after controlling for the bid-ask spread, trading these portfolios has been significantly profitable during the sample period from 1996 to 2000.

O'Brien and Shackleton (2005) examine the effect of systematic moments of order higher than two in explaining the cross-section of option returns. They focus on FTSE100 index options and conclude that while systematic variance is significant in explaining option returns, the effect of coskewness and cokurtosis is less evident. Finally, one of the papers that have significantly motivated the present study is the Santa-Clara and Saretto (2009) examination of S&P 500 options returns. In particular, they analyze the performance of various trading strategies and find that these strategies are associated with very high returns, an effect that is especially pronounced for those that involve short positions in options, and that these returns are not justified by their exposure to market risk according to traditional asset pricing models. However, after accounting for transaction costs and margin requirements, the above mentioned returns become less significant, or even negative. Consequently, even though a certain level of mispricing is documented in the US options market, a typical investor cannot exploit real profit opportunities due to the high costs involved and, instead of being arbitraged away, options mispricing is allowed to persist.

Most empirical studies that examine options and other derivatives instruments have traditionally focused on developed markets such as the US and the UK option markets. Although this is partly justified by the fact that such markets are characterized by high-volume trading so that option prices are likely to be more informative, this Chapter proposes that emerging markets can provide an interesting new field of research.

The Athens Derivatives Exchange (ADEX) in particular was established in 2000 and has experienced significant growth since, being ranked 7th among European option exchanges based on volume in index derivatives in 2003.[9] The relatively limited research interest in the Greek derivatives market has so far focused mainly on futures rather than on options. For instance, Floros and Vougas (2006) examine the hedging effectiveness of Greek stock index futures, while Floros (2007) and Kenourgios (2004) explore the price discovery mechanism that links index futures contracts with the underlying index. Skiadopoulos (2004) is the only paper to date that focuses on the Greek options market and, in particular, implied volatility. Using a dataset of options on the Greek large-capitalization FTSE/ASE-20 index, the author constructs an implied volatility index GVIX and finds that, although the spot level of the underlying index can forecast future changes in implied volatility, the reverse is not true, supporting the hypothesis that '... GVIX can be interpreted as a gauge of the investor's sentiment...' rather than a predictor of future realized volatility. Furthermore, Skiadopoulos (2004) documents a spillover effect between the GVIX and the US volatility indices VXO and VXN.

[9] Athens Derivatives Exchange Fact Book (2005)

5.1.2 Scope of Study

The objective of this Chapter is to examine option returns in an emerging market with particular focus on the extent of mispricing present in the market. In addition, a comparison is provided between the degree of efficiency exhibited by the emerging ADEX and that which typically characterizes developed options markets, with the main focus of the comparison being the US. However, the above comparison is not a direct, one-to-one comparison between exchanges. More specifically, the efficiency of the Greek market is evaluated by performing CAPM regressions on option returns and by examining the returns of delta and/or vega neutral straddles in ADEX for the period 2004-2007. Rather than replicating this analysis for the US market in the same period, empirical findings from previous studies which employ similar efficiency tests in developed markets across much larger sample periods are used as a benchmark.

The main hypothesis of interest is that the Athens Derivatives Exchange exhibits a level of efficiency comparable to that of developed markets. Given the global nature of today's marketplace and the fact that large, international investors with significant experience in more established options markets account for most of the trading volume in Greece, it is not unreasonable to assume that ADEX should be a relatively efficient market, with option prices reflecting 'true' asset values that do not offer returns in excess of those justified by their risk-exposure.

The alternative hypothesis is partly motivated by the Santa-Clara and Saretto (2009) study, and states that higher transaction costs, combined with thinner trading, are likely to be associated with a higher level of options mispricing, measured by returns of options and options strategies in excess of their exposure to risk. Intuitively, higher trading costs will result in a widened no-arbitrage band, and market prices of options will be allowed to deviate further from their theoretical price without arbitrageurs being able to profit from the discrepancy and, in the process, forcing prices to their 'true' level.

Given the higher transaction costs that characterize the Athens Derivatives Exchange, the above mentioned alternative hypothesis predicts that positions in individual options in Greece earn higher risk-adjusted returns than those typically earned by options in the US or the UK developed markets. In addition, trading strategies that are risk-neutral would be more likely to earn returns that are statistically different from the risk-free rate in an emerging market compared to its developed counterparts.

Overall, the results appear to support the efficiency of the Athens Derivatives Exchange. Although naked option positions in Greece earn substantially higher returns than their US counterparts, the discrepancy between realized returns and those justified by standard asset pricing theory or the Black and Scholes option pricing model (1973) is not necessarily larger than that traditionally documented in the US market. More importantly, portfolios that are formed to be delta and/or vega neutral are found to earn the risk-free rate, providing further support for the efficiency of the Greek options market. In summary, the developing market of the Athens Derivatives Exchange does not appear to offer real profit opportunities, after controlling for risk, and the extent to which options might be considered as mispriced is not found to be higher than that characterizing the US and the UK markets.

The remaining of the Chapter is organized as follows. Section 5.2 gives an overview of the Greek large-capitalization index and its returns throughout the sample period, while Section 5.3 presents the data used in the empirical analysis. Section 5.4 describes the estimation of the implied moments vector for the index using options data, and Section 5.5 discusses observed returns of naked positions in individual, European-style calls and puts written on the FTSE/ASE-20. Section 5.6 analyzing risk-adjusted returns. Section 5.7 includes the analysis of returns to various trading strategies, such as delta and vega neutral straddles. Finally, Section 5.8 provides a comparison between the results and previous empirical findings from developed markets, and Section 5.9 concludes.

5.2 Index Returns

The sample period could be characterized as one of a significantly high increase in the level of the underlying large-capitalization FTSE/ASE-20 index. From a level of 1,194.2 on the first trading day of 2004, the index has experienced a rapid growth to reach 2,566 at the end of January 2007. This translates into an overall appreciation of 114.87% over the entire 37 months period or, equivalently, 28.15% annually. The evolution of the FTSE/ASE-20 throughout the sample period is presented in Figure 5.1.

Daily returns are plotted in Figure 5.2. The mean daily return is 0.1% with a standard deviation of 105 basis points and, as can be seen from Figure 5.3, the returns' distribution seems to deviate from normality.[10] In addition, spot returns appear to be an I(0) process, since the Dickey-Fuller test for the order of integration produces a t-statistic equal to 1.64 and therefore cannot reject the null hypothesis of a unit root in the time-series of arithmetic returns at the 5% significance level.

Figure 5.1

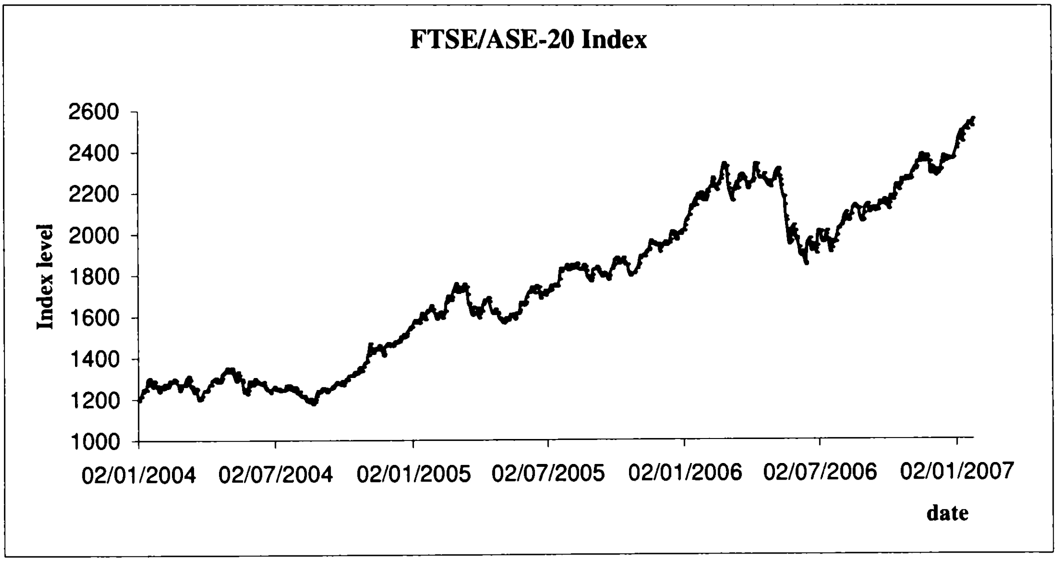

[10] Furthermore, the Jarque-Bera test rejects the null hypothesis for normality in returns and log-returns at the 5% significance level.

Figure 5.2

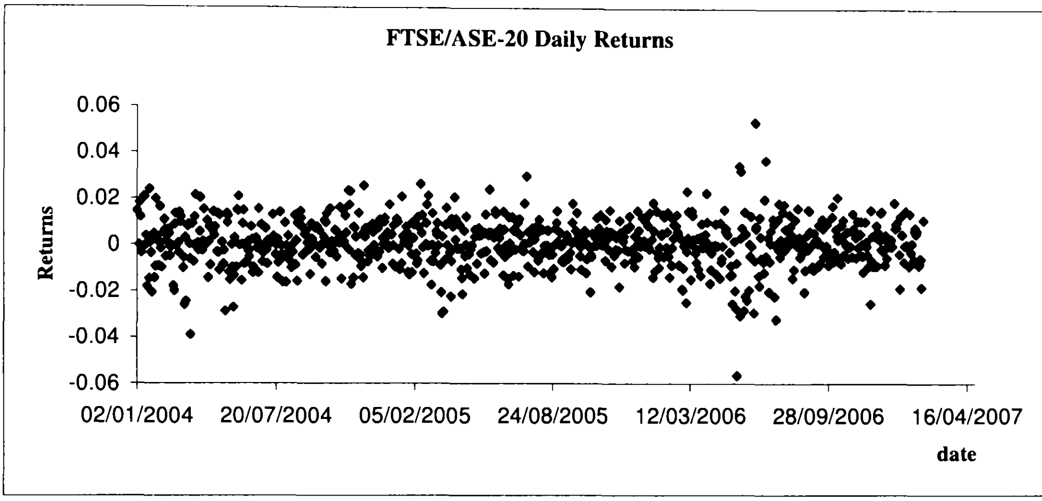

Figure 5.3

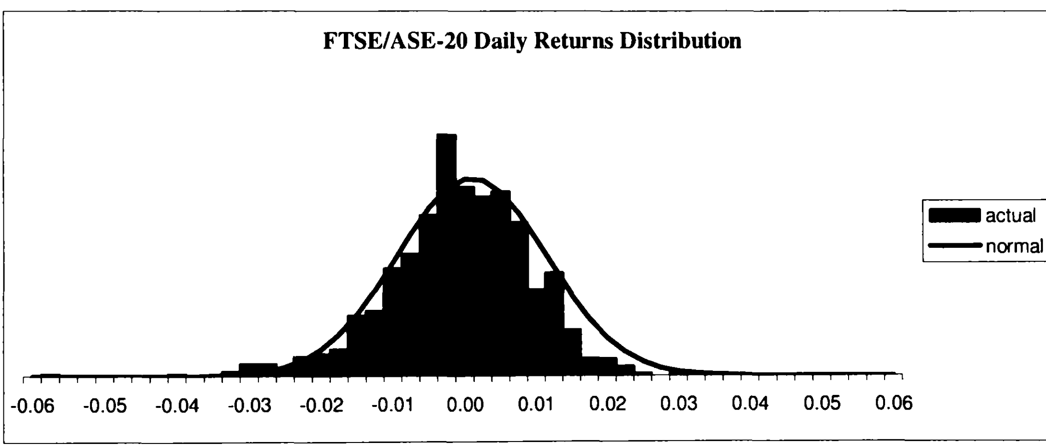

Not surprisingly, returns of the FTSE/ASE-20 index exhibit an extremely high correlation with market returns, which are proxied by the Athens Composite Share Price Index (ACSPI). In addition to a correlation coefficient of 0.98, this is farther supported by the estimated betas of the index within the standard CAPM framework. More specifically, rolling 120-days regressions of the CAPM equation described in (5.1) are estimated,

$$E[R_{in}] = \alpha + \beta E[R_m] + \varepsilon \tag{5.1}$$

where R_{in} is the excess return of the FTSE/ASE-20, R_m is the market risk-premium and ε is a random error term. When stated in excess returns, the null hypothesis of the CAPM is that $\alpha = 0$. As can be seen from Figure 5.4, the resulting betas are very close to unity throughout the entire sample period, reflecting the fact that the 20 largest capitalization stocks that are included in the FTSE/ASE-20 heavily influence the overall market index.

Figure 5.4

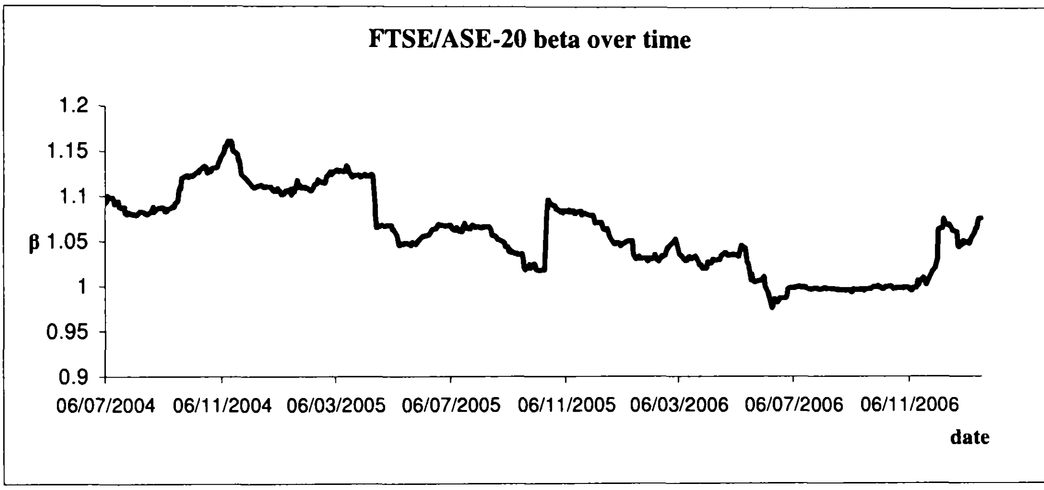

5.3 Data

The original dataset consisted of 15,198 calls and 18,217 puts traded on the Derivatives Market of the Athens Stock Exchange. The options are European-style and written on the FTSE/ASE-20 index which includes the 20 most liquid and largest capitalization Greek stocks. All relevant options data is publicly available through the exchange's website (www.adex.ase.gr). For every calendar day, option prices are obtained for the two nearest expiration dates which are typically more liquid than longer-term contracts. The options expire on the third Friday of the month and settlement is in cash.

The dataset runs from January 2004 to January 2007 for a total of 770 trading days. Similarly to previous studies, several filters are employed. First, all options with prices that lay outside the well-known theoretical bounds or are near zero are excluded from the dataset. Moreover, calls and puts with less than one week (five trading days) to maturity are dropped. Finally, options with less than five traded contracts on a given day are excluded to avoid illiquidity concerns. The above filters resulted in a reduced dataset of 9,761 calls and 9,212 puts.

Table 5.1 reports some descriptive statistics with respect to the number of daily observations. Although the sample consists of 12.34 calls and 11.90 puts on average per calendar day, it can be seen from Figure 5.5 that the number of observations in a day exhibits significant variability throughout the period of January 2004 to January 2007. The number of calls (puts) per day exhibits a standard deviation of 3.26 (2.99) and ranges from a minimum of 5 to a maximum of 23 (24).

Table 5.1

Option Observations per Calendar Day

	Calls	Puts
Mean	12.34	11.90
Median	12	12
Standard deviation	3.26	2.99
Minimum	5	5
Maximum	23	24

This Table presents descriptive statistics of the options dataset used. Calls and puts are tabulated separately.

The risk-free interest rate is proxied by Euribor which, along with the underlying FTSE/ASE-20 index, was obtained through DataStream. The Athens Composite Share Price Index was obtained by Reuters. The dividend yield of the underlying asset was calculated by using futures contracts on the FTSE/ASE-20 index (futures data are also available through the exchange's website) and solving equation (5.2) for the dividend yield:

$$F_0^T = S_0 e^{(r-q)T} \tag{5.2}$$

where F_0^T is the value at time 0 of a futures contract on the index expiring at T, S_0 is the spot price of the index, r is the risk-free rate and q is the dividend yield.

Figure 5.5

Option Observations per Trading Day

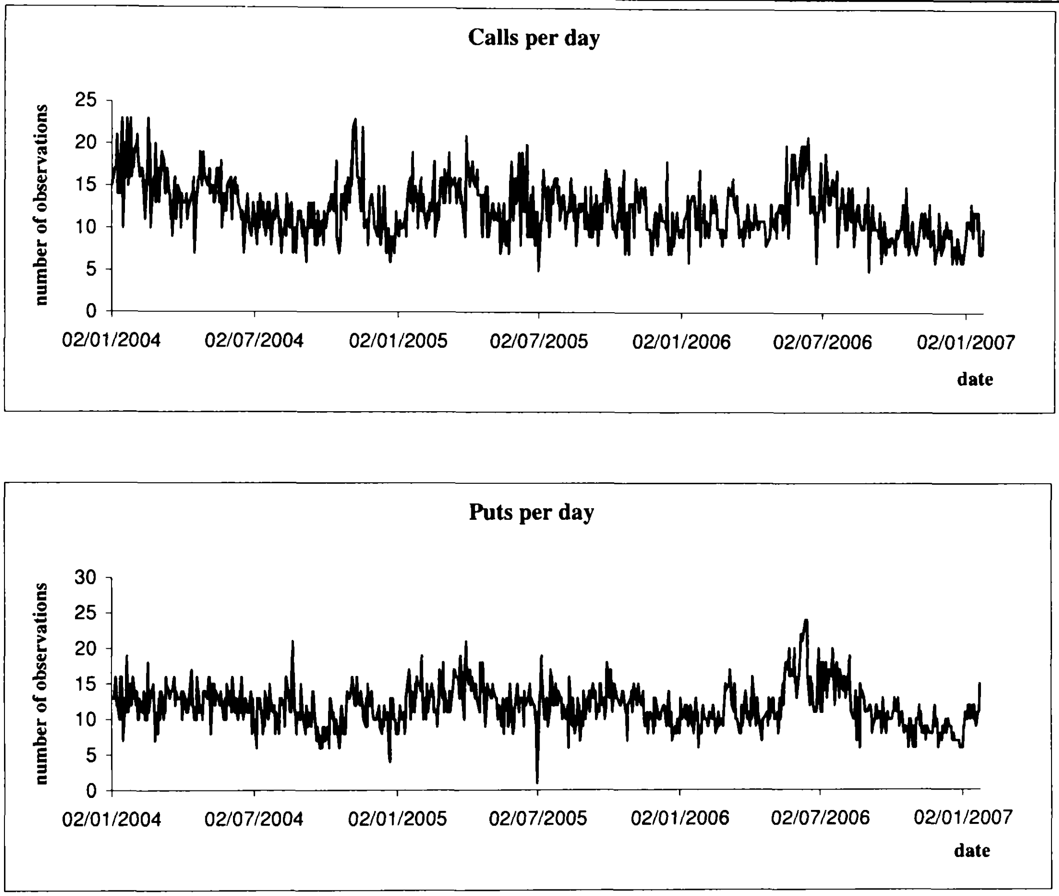

Before proceeding with the empirical analysis, the limitation of using a relatively small sample should be acknowledged. The Athens Derivatives Exchange is a substantially new market that was established in 2000 and, consequently, trading volume remained at fairly low levels during its first operating years. Trading activity picked up significantly, though, after 2004 and this essentially introduces a lower bound to the sample period that can be examined. As has already been mentioned in Section 5.2, the sample period of January 2004 to January 2007 corresponds to a significantly 'bull' market in Greece with the underlying large-capitalization FTSE/ASE-20 index appreciating by more than 28% per year. Therefore, the fact that such a substantially high risk-premium characterizes a relatively small sample period will undoubtedly have an effect on the sign as well as on the magnitude of observed option returns in ADEX.

More specifically, Coval and Shumway (2001) demonstrate that under the moderate assumptions of the stochastic discount factor being negatively correlated with the underlying and of the underlying having a positive expected return (which is obvious in the case of a market index in the long term), any call option written on the underlying *must* have a positive expected return that is higher than the underlying's expected return, and any put option on the underlying *must* have an expected return below the risk-free rate (see Propositions 1 and 2, respectively, in Coval and Shumway (2001)). Given the high positive market premium in Greece during 2004-2007, it is extremely likely that call options, which represent levered positions in the underlying, will offer returns that are positive and significantly high (at least as high as the 28% market premium and dependent on the option's moneyness, i.e. its leverage), while put returns will be negative and of high absolute magnitude.[11]

5.4 Implied Moments

5.4.1 Theoretical framework

This Section describes the estimation of the underlying distribution that is implied by option prices. A large part of past studies has used the Black and Scholes (1973) option pricing formula to infer implied volatility. Although this has been a significantly popular methodology throughout the related literature, the Black and Scholes (B&S) implied second moment relies on certain assumptions, the most restrictive of which is arguably the log-normality of the asset's distribution. Under this set of assumptions, the theoretical price C_{BS} of a European call option is given by:

$$C_{BS} = S_0 e^{-qT} N(d_1) - K e^{-rT} N(d_2) \tag{5.3}$$

[11] See Appendix B for a more detailed discussion of the Coval and Shumway (2001) predictions on theoretical expected option returns.

with
$$d_1 = \frac{\ln(\frac{S_0}{K}) + (r - q + \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}$$

$$d_2 = d_1 - \sigma\sqrt{T}$$

where S_0 is the price of the underlying asset, K is the call's exercise price, r is the constant risk-free interest rate, q is the asset's constant dividend yield, T is the time-to-maturity, and $N(\cdot)$ is the standard normal cumulative distribution function.

Skiadopoulos (2004) computes an implied volatility index GVIX for the FTSE/ASE-20 based on the B&S formula. Using options and futures written on the index for the last three months of 2002, he estimates the B&S implied volatility of the first OTM call and that of the first OTM put. The GVIX implied volatility index is then defined as the linear interpolation between these two moments, i.e. the implied volatility of the above OTM synthetic option.

This Section intends to improve on Skiadopoulos (2004) in two ways. First, a much larger sample is used in order to observe the variation across time of implied volatility. Most importantly, though, an alternative option pricing model is adopted, namely the Corrado & Su (1996) option pricing formula, in an attempt to account for the observed deviation of returns from log-normality. This framework also has the advantage of simultaneously estimating the next two implied moments of the asset's distribution in addition to implied volatility.

The Corrado and Su (1996) methodology relaxes the log-normality restriction by modifying the original B&S formula to account for non-zero skewness and excess kurtosis in the underlying's distribution using a Gram-Charlier series expansion of the standard normal density function. More specifically, they define a density function $g(z)$ which accounts for non-normal skewness and kurtosis, described by the following equation, where $n(z)$ represents the standard normal density function and μ_n is the standardized coefficient of the nth moment of the asset's returns distribution.

$$g(z) = n(z)[1 + \frac{\mu_3}{3!}(z^3 - 3z) + \frac{\mu_4 - 3}{4!}(z^4 - 6z^2 + 3)] \tag{5.4}$$

Within this framework, a call option's price is expressed as the sum of the theoretical B&S price and two correction terms related to non-normal skewness and excess kurtosis. Then, the Corrado and Su (C&S) call option price C_{CS} is given by:

$$C_{CS} = C_{BS} + \mu_3 Q_3 + (\mu_4 - 3) Q_4 \qquad (5.5)$$

with

$$Q_3 = \frac{1}{3!} S_0 \sigma \sqrt{T}[(2\sigma\sqrt{T} - d_1)n(d_1) + \sigma^2 T N(d_1)]$$

$$Q_4 = \frac{1}{4!} S_0 \sigma \sqrt{T}[(d_1^2 - 1 - 3\sigma\sqrt{T}(d_1 - \sigma\sqrt{T}))n(d_1) + \sigma^3 T^{3/2} N(d_1)]$$

where C_{BS} is the theoretical Black and Scholes option price, Q_3 and Q_4 represent the marginal effects of non-normal skewness and kurtosis for the option's price, respectively, and $n(\cdot)$ is the standard normal density function.

5.4.2 Empirical Estimation

For each calendar day, option observations are separated based on their respective maturities. Obviously, the first options set includes calls that expire on the nearest expiration date, while the second set includes calls expiring on the second-nearest expiration date. The differences between market prices C_M and theoretical C&S call prices C_{CS} are then computed for each of the above two groups. The implied moments are estimated as the values of the moment vector (σ, μ_3, μ_4) that minimize the sum of the following squared errors:

$$\min_{(\sigma^T, \mu_3^T, \mu_4^T)} = \sum_{i=1}^{K} [(C_M^T(S, K_i) - C_{CS}^T(S, K_i))^2]$$

$$\min_{(\sigma^{T'}, \mu_3^{T'}, \mu_4^{T'})} = \sum_{i=1}^{K'} [(C_M^{T'}(S, K_i) - C_{CS}^{T'}(S, K_i))^2]$$

where T and T' refer to the nearest and second-nearest expirations, respectively. K and K' denote the number of observations for T and T', respectively. This procedure provides two implied vectors for every calendar day, corresponding to different expirations. In order to standardize the results to a constant time-horizon, a linear interpolation is used between times T and T', as shown by equations (5.6), (5.7) and (5.8), to estimate the 30-day period implied volatility, skewness and kurtosis, respectively.

$$\sigma^{30} = (\frac{T'-30}{T'-T})\sigma^{T} + (\frac{30-T}{T'-T})\sigma^{T'} \tag{5.6}$$

$$\mu_3^{30} = (\frac{T'-30}{T'-T})\mu_3^{T} + (\frac{30-T}{T'-T})\mu_3^{T'} \tag{5.7}$$

$$\mu_4^{30} = (\frac{T'-30}{T'-T})\mu_4^{T} + (\frac{30-T}{T'-T})\mu_4^{T'} \tag{5.8}$$

For notational purposes, the time indicators are dropped throughout the subsequent analysis, with σ, μ_3 *and* μ_4*-3* denoting 30-day standardized estimates of implied volatility, skewness and excess kurtosis, respectively, for the 30-day horizon. Figures 5.6 to 5.8 present the estimated vectors of implied moments for the above two option groups across time, while Figures 5.9 to 5.11 plot the standardized estimates.

Figure 5.6

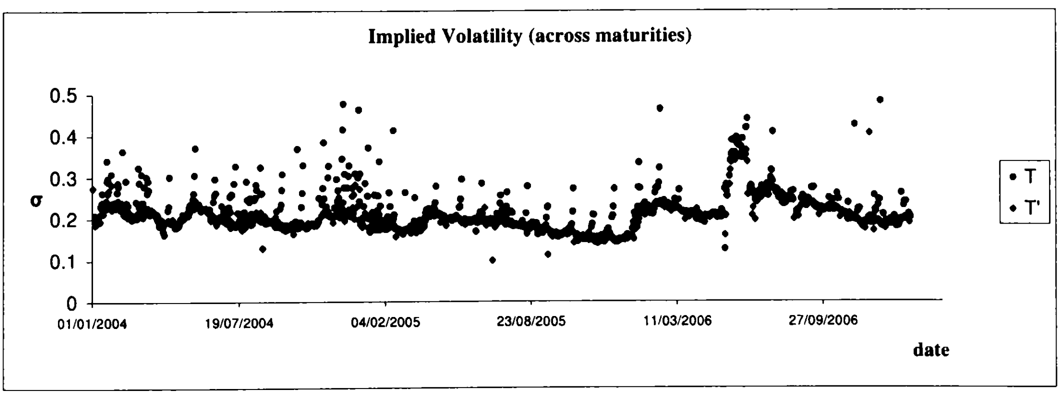

Figure 5.7

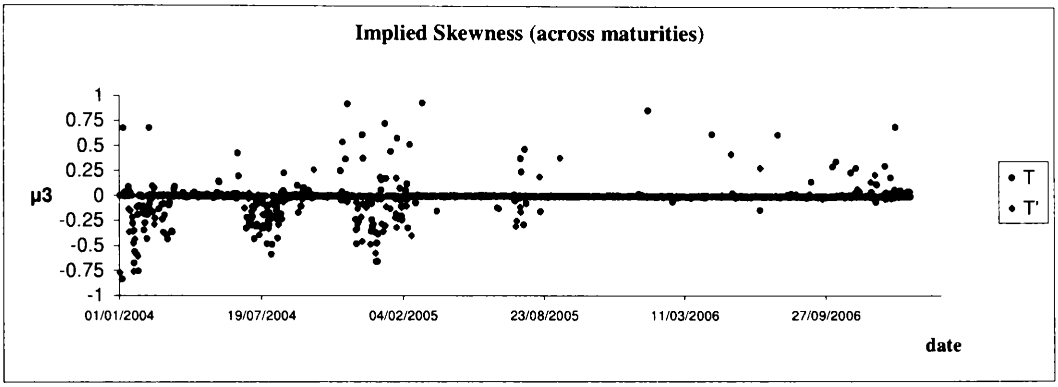

Figure 5.8

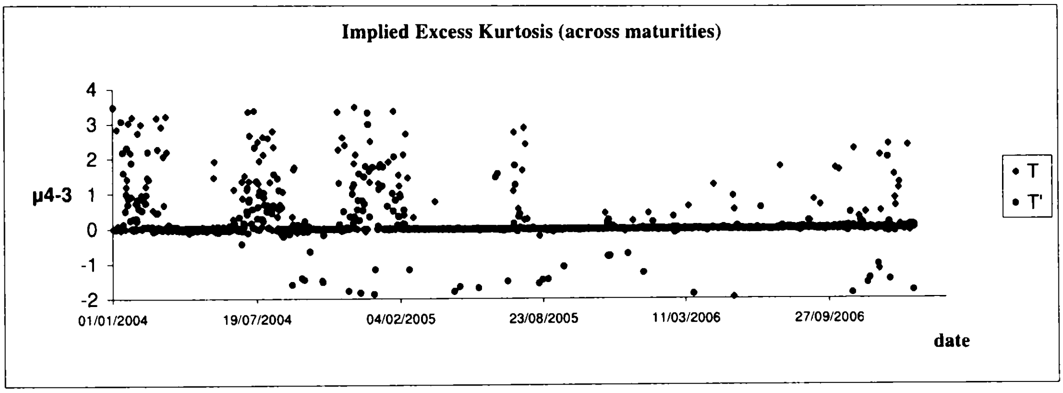

Figure 5.9

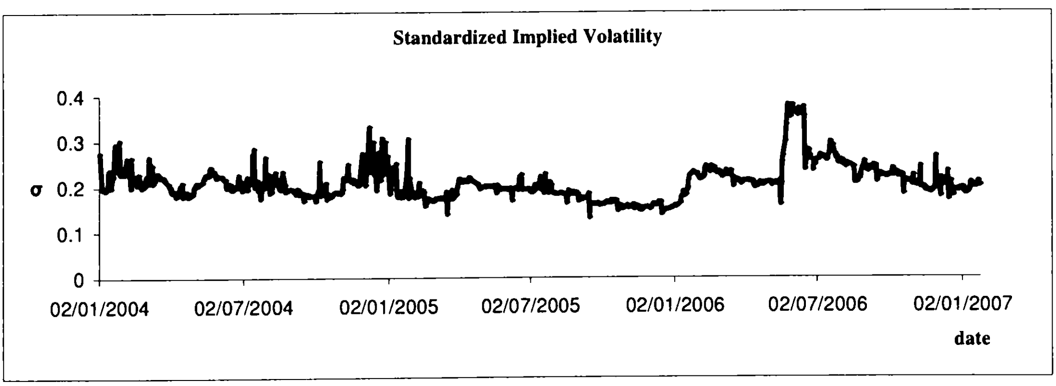

Figure 5.10

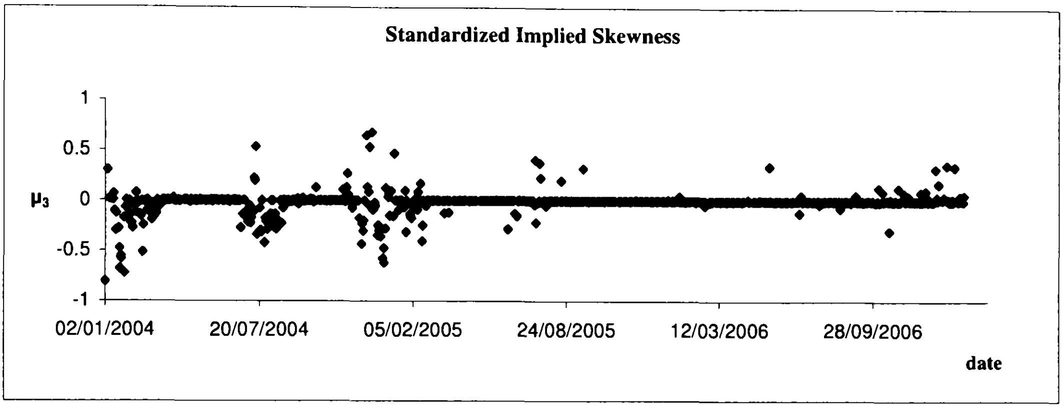

Figure 5.11

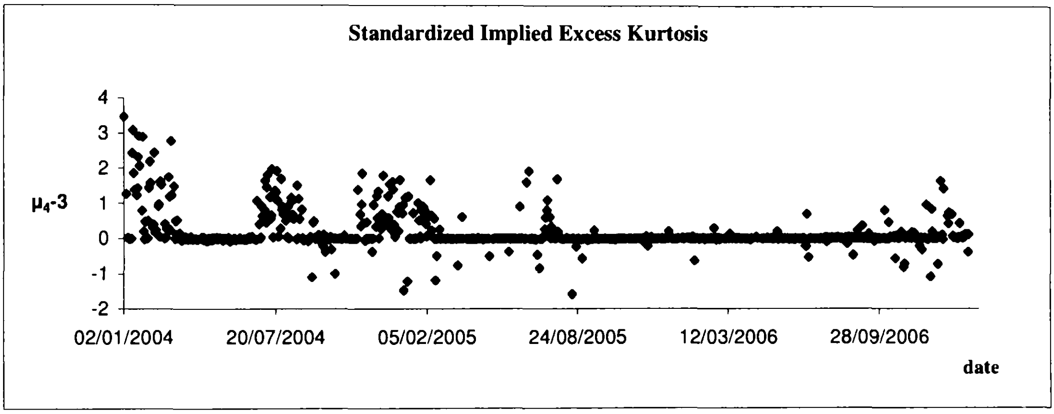

The implied volatility σ has a mean (median) value of 20.93% (20.36%) and ranges between 13.19% and 37.89%. Periods of high and low volatility are easily observable from Figure 5.9, indicating the presence of some autocorrelation in the time-series. Although implied skewness is statistically indistinguishable from zero for most of the sample period, it can be seen from Figure 5.10 that there are certain periods of relatively large positive or negative observations. The most easily identifiable such periods include the first two months of 2004, which are characterized by negative skewness with a minimum of -0.8, as well as the period from July to mid August 2004 which also exhibits negative skewness. Finally, from November 2004 to February 2005, the implied third moment experiences its most volatile period with many extreme observations, both positive and negative. Implied excess kurtosis exhibits a somewhat similar pattern. Despite the fact that it is near zero in most sample days, the above mentioned periods of extreme skewness are also characterized by non-normal kurtosis. The main difference in the two time-series is the fact that extreme excess kurtosis does not change signs as skewness does, with most outliers lying above zero.

Table 5.2

Implied Moments

Panel A: Implied Moments' Descriptive Statistics			
	Volatility σ^{30}	Skewness μ_3^{30}	Excess Kurtosis $(\mu_4\text{-}3)^{30}$
Mean	0.209	-0.018	0.178
Median	0.204	-0.001	0.009
St. Deviation	0.0383	0.117	0.518
Minimum	0.132	-0.806	-1.59
Maximum	0.379	0.678	3.472
Panel B: Jarque-Bera Test for Normality			
	Volatility σ^{30}	Skewness μ_3^{30}	Excess Kurtosis $(\mu_4\text{-}3)^{30}$
Null	Reject	Reject	Reject
p-value	0	0	0
t-statistic	870.27	6,591	3,249
Critical value	5.99	5.99	5.99
Panel C: Order of Integration			
	Volatility σ^{30}	Skewness μ_3^{30}	Excess Kurtosis $(\mu_4\text{-}3)^{30}$
Dickey-Fuller I(0) t-stat	17.12	13.36	14.72
Dickey-Fuller I(1) t-stat	-136.81	-18.31	-17.49
GPH (d)	0.59	0.27	0.45

This Table presents an analysis of the FTSE/ASE-20 implied moments for the period January 2004 to January 2007. Panel A provides descriptive statistics, Panel B tabulates the results of the Jarque-Bera normality test and Panel C discusses the order of integration of the implied moments based on the Dickey-Fuller test as well as on the Geweke and Porter-Hudak estimator.

Panel A of Table 5.2 presents descriptive statistics for the time-series of the three implied moments. Panel B reports the results of the Jarque-Bera test which rejects normality for all variables of interest, while Panel C provides the results for the Dickey-Fuller and the Geweke & Porter-Hudak (GPH) test for the variables' order of integration. As can be seen from the Table, the Dickey-Fuller test rejects the null hypotheses of σ

being either an I(0) or an I(1) process while the GPH estimator suggests that implied volatility is a fractionally integrated process I(d) with d equal to 0.59. Furthermore, similarly to the implied volatility series, the two higher implied moments exhibit characteristics of fractional integration. More specifically, the Dickey-Fuller test rejects both null hypotheses of I(0) and I(1) for skewness and kurtosis at the 5% significance level, while the GPH test estimates skewness and excess kurtosis to be fractionally integrated processes I(d) with d equal to 0.27 and 0.45, respectively.

5.4.3 Alternative Specification

In addition to the above methodology for estimating the implied vector [σ, μ_3, μ_4], an alternative estimation of implied moments is performed. The main difference is the fact that option observations on a given day are not grouped according to their respective maturities, as was the case in Section 5.4.2. Instead, the previously mentioned sum of squared errors between market prices for calls C_M and theoretical Corrado and Su prices C_{CS} is minimized across the entire set of options per day, irrespective of their time to expiration.

$$\min SSE(\sigma^{ALT}, \mu_3^{ALT}, \mu_4^{ALT}) = \sum_{K=1}^{K} [(C_M(S,K) - C_{CS}(S,K))^2]$$

Figure 5.12 plots the estimated implied moments that resulted from this alternative methodology. When compared to the estimates in Figures 5.9 to 5.11, it can be easily seen that, although some patterns are present in both sets of time-series, there are also significant differences since all implied moments exhibit higher volatility under the alternative estimation.

Figure 5.12

Implied Moments from Alternative Methodology

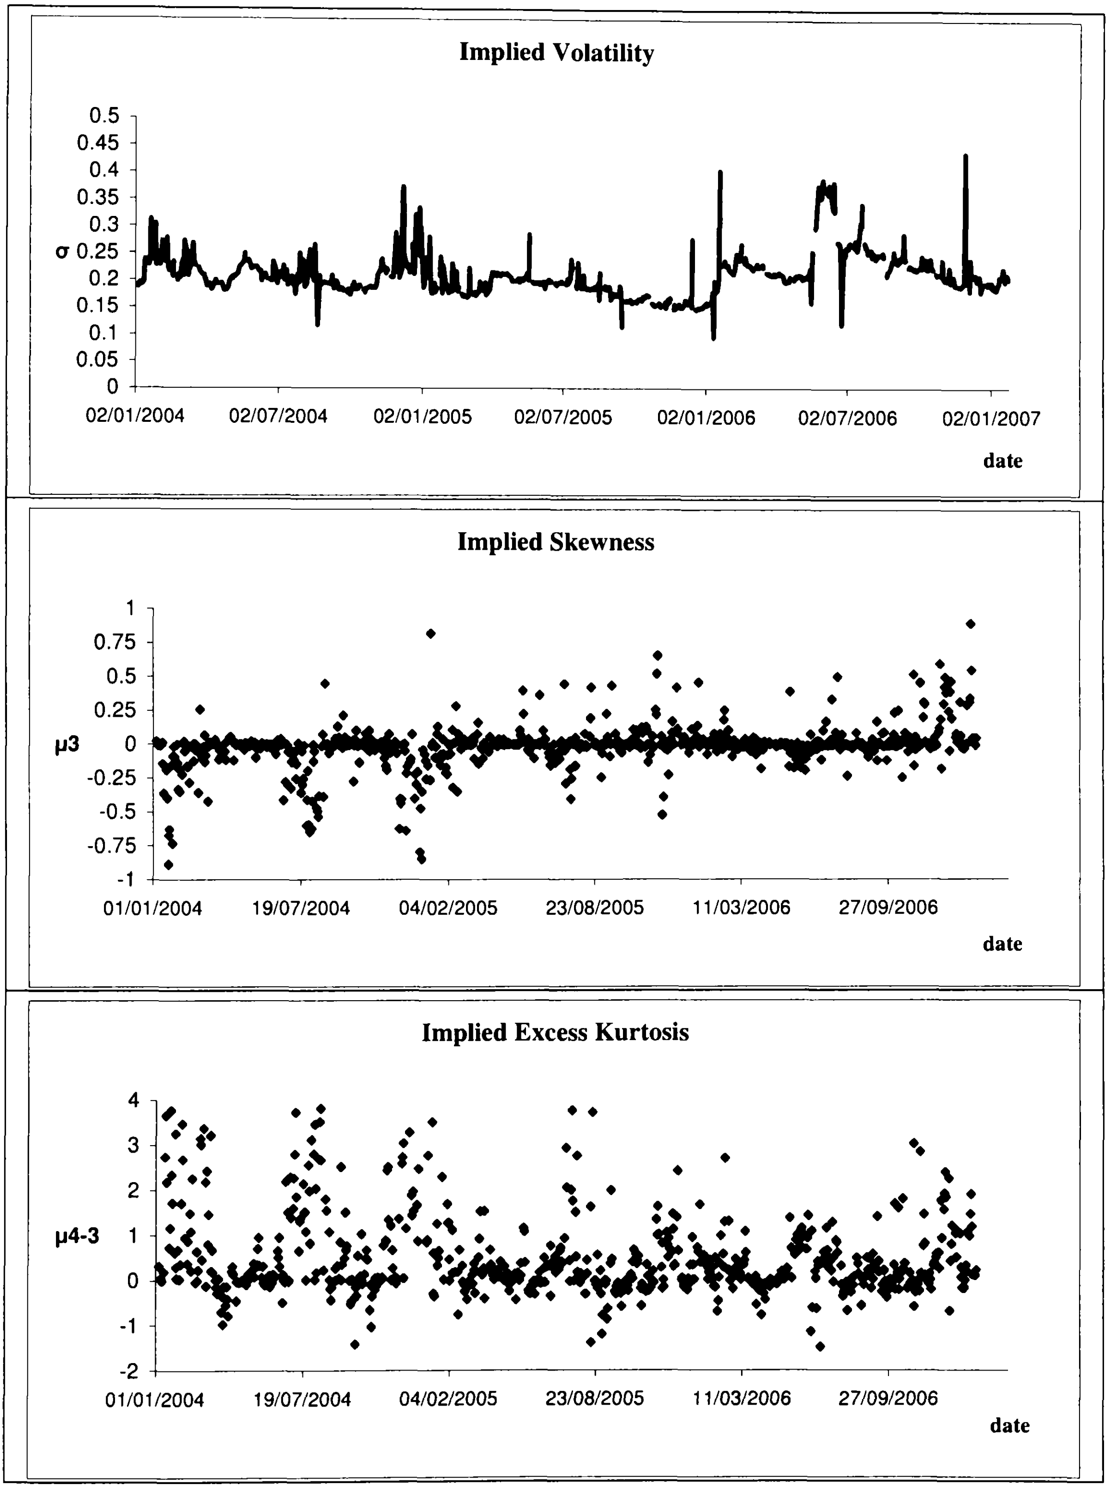

The main disadvantage of this alternative procedure is that it provides only one vector of implied moments. Since options with different expirations are combined in the minimization of the SSE, the time horizon that corresponds to the estimated vector is not clear and the resulting SSE are significantly higher than those in Section 5.4.2. More specifically, the average SSE for the first methodology, weighted by the number of available options per expiration date, is 1.22, compared to an average of 7.76 for the second one. The time series of SSE for both estimations are plotted in Figure 5.13.

Figure 5.13

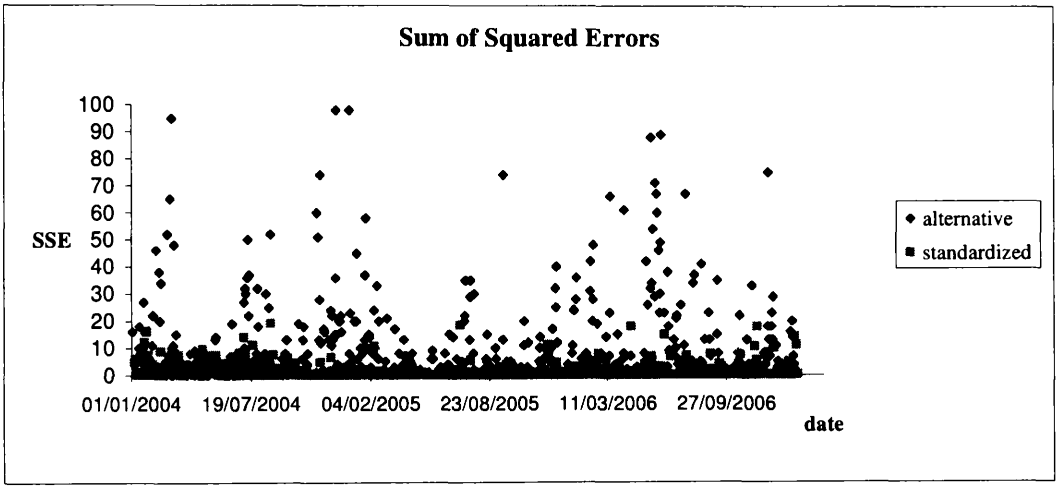

5.5 Option Returns

5.5.1 Call Options

As has been already mentioned in Section 5.3, the call options dataset (post filtering) consists of 9,761 call observations across 770 trading days. Daily arithmetic returns for each individual call are computed using closing prices for each calendar day. Obviously, the fact that not all strike prices have traded options on every day reduces the number of calls for which daily returns can be computed. Whenever a specific call is not traded on two consecutive trading days, or the call price remains the same over this window, the call return is treated as a missing observation. The above limitations reduce the number of computable daily returns to 6,884 observations.

Let c_t be the price of a call option with strike price K and time-to-maturity T. The daily arithmetic return R_c for this call is estimated as the difference between c_{t+1} and c_t, divided by c_t.

$$R_c = \frac{c_{t+1} - c_t}{c_t} \tag{5.9}$$

Options with different strike prices are likely to earn returns that differ significantly. For instance, Coval and Shumway (2001) show that option returns should be increasing across strike price space. In order to examine the behaviour of call returns across different strikes, individual calls are sorted into four groups according to their moneyness, using three different moneyness proxies (see also Ni (2006)). These moneyness groups are created such that strike prices for a given call are increasing across the strike group number, with group 1 including calls with the lowest strikes and group 4 including calls with the highest strikes.

The first moneyness proxy refers to the ratio of the strike price to the price of the underlying index. Being the most typical way of defining moneyness, this ratio ensures that different strike prices are normalized by the closing level of the FTSE/ASE-20. The second way of classifying option returns is to divide the logarithm of the above ratio with the volatility of the underlying. For each calendar day, the historical volatility σ_t of the index is estimated over the previous 60 trading days:

$$\sigma_t = \sqrt{\frac{250}{60} \sum_{i=0}^{59} R_{ind,t-i}^2} \tag{5.10}$$

where $R_{ind,t-i}$ is the return of the FTSE/ASE-20 index on trading day t-i. This method has the additional advantage of controlling for the underlying's volatility. The third moneyness proxy is the Black and Scholes option's delta, with equation (5.10) used to estimate volatility in the B&S formula. The B&S delta simultaneously accounts for differences in underlying index level, index volatility and time-to-maturity across individual options. Panel A of Table 5.3 contains the criteria for assigning call option returns to the four strike groups.

After classifying calls according to the above mentioned criteria, summary statistics for call returns of each of the strike groups are computed. Table 5.4 presents the mean, standard error and skewness of call returns across the four groups. T-statistics for the null hypothesis that the average call return is statistically indistinguishable from zero are in brackets. Finally, the average call B&S beta and average volume of traded contracts for each option category are reported. The B&S beta β_{BS} for each call is estimated using equation (5.11) and is of particular interest, since standard asset pricing theory predicts that average call returns should increase as β_{BS} increases. Following Coval and Shumway (2001), this also implies that β_{BS} will be higher for calls with higher strikes than for their lower strikes counterparts. The intuition behind this theoretical prediction is that calls with higher strike prices represent more levered positions in the underlying asset and are, therefore, riskier investments.

Table 5.3

Strike Groups

Panel A: Call Options				
Strike Group	1	2	3	4
R = K/S	$R \leq 0.95$	$0.95 < R \leq 1$	$1 < R \leq 1.05$	$1.05 < R$
R = ln(K/S)/σ	$R \leq -0.05$	$-0.05 < R \leq 0$	$0 < R \leq 0.05$	$0.05 < R$
R = B&S delta	$0.75 \leq R \leq 1$	$0.5 \leq R < 0.75$	$0.25 \leq R < 0.5$	$0 \leq R < 0.25$
Panel B: Put Options				
Strike Group	1	2	3	4
R = K/S	$R \leq 0.98$	$0.98 < R \leq 1$	$1 < R \leq 1.02$	$1.02 < R$
R = ln(K/S)/σ	$R \leq -0.20$	$-0.20 < R \leq 0$	$0 < R \leq 0.20$	$0.20 < R$
R = B&S delta	$-0.25 \leq R \leq 0$	$-0.5 \leq R < -0.25$	$-0.75 \leq R < -0.5$	$-1 \leq R < -0.75$
Panel C: Delta-Neutral Straddle Portfolios				
Strike Group	1	2	3	4
R = K/S	$R \leq 0.98$	$0.98 < R \leq 1$	$1 < R \leq 1.02$	$1.02 < R$
R = ln(K/S)/σ	$R \leq -0.20$	$-0.20 < R \leq 0$	$0 < R \leq 0.20$	$0.20 < R$
R = B&S delta	$-1 \leq R \leq -0.50$	$-0.50 < R < 0$	$0 \leq R < 0.50$	$0.50 \leq R < 1$

This Table presents the cutoff points for assigning calls, puts and delta-neutral straddles into moneyness groups.

$$\beta_c = \frac{S}{c} N[\frac{\ln(\frac{S}{X}) + (r - q + \frac{\sigma^2}{2})t}{\sigma\sqrt{t}}]\beta_S \tag{5.11}$$

As can be seen from Table 5.4, average daily arithmetic returns for call options in the Greek market are positive and particularly high, compared, for instance, to call returns in the US. Panel A refers to strike groups according to the strike to underlying ratio, Panel B to the logarithm of strike-to-underlying divided by σ, and Panel C to the B&S delta moneyness proxy. For the first classification method, call returns are found to be statistically significant from groups 2 to 4, and marginally significant for the first group, ranging from a minimum of 1.15% daily for the low-strike, most in-the-money calls to a maximum of 3.81% for the high-strike, most out-of-the-money ones. These figures correspond to annual returns of roughly between 288% and 953%, depending on the options' moneyness, and are much higher than returns of calls written on the S&P 500, which have been around 100% per annum (see Coval and Shumway (2001)). Results from Panels B and C are similar, with the exception of call returns for the highest-strike group in Panel C being only marginally significant, although relatively high in absolute magnitude.

Furthermore, average call returns appear to support theoretical predictions, in the sense that they are strictly increasing across strike price space. The average B&S betas for the four groups are also increasing as the strike increases, indicating that options which exhibit a higher correlation with the market tend to earn higher returns on average than options which are more weakly correlated with the market. For instance, when calls are assigned to *K/S* groups, β_{BS} ranges from 14.07 for the low-strike group to 36.71 for the high-strike one.

Table 5.4

Summary Statistics for Call Options and their Daily Returns (closing price)

Strike Group	1	2	3	4
		Panel A: K/S		
Call Beta	14.07	22.70	32.29	36.71
Call Volume (Contract)	29.09	108.16	209.73	118.32
Average Call Return	0.0115	0.0222	0.0242	0.0381
t-stat	(1.68)	(4.13)	(3.54)	(2.53)
Call Return St. Dev	0.1644	0.2601	0.3725	0.4809
Call Return Skewness	0.0779	0.2988	0.8601	1.4085
No. of Obs	576	2,329	2,963	1,014
		Panel B: ln(K/S)/σ		
Call Beta	14.23	22.76	32.33	36.78
Call Volume (Contract)	29.06	109.06	208.37	118.48
Average Call Return	0.0130	0.0220	0.0240	0.0393
t-stat	(1.90)	(4.05)	(3.53)	(2.54)
Call Return St. Dev	0.1677	0.2605	0.3738	0.4819
Call Return Skewness	0.1355	0.2987	0.8537	1.4381
No. of Obs	602	2,303	3,006	971
		Panel C: B&S delta		
Call Beta	15.87	21.73	30.90	44.42
Call Volume (Contract)	41.31	99.81	195.30	135.59
Average Call Return	0.0141	0.0159	0.0299	0.0334
t-stat	(1.90)	(3.07)	(4.95)	(1.50)
Call Return St. Dev	0.1721	0.2438	0.3540	0.5866
Call Return Skewness	-0.1030	0.2828	0.8311	1.1761
No. of Obs	538	2,195	3,453	697

This Table tabulates summary statistics of call returns across the period January 2004 to January 2007. Calls have been assigned to moneyness groups according to the cutoff points in Panel A of Table 5.3.

Another interesting finding is the monotonic relationship between the variability and the skewness of call returns, and the strike group. Call returns are found to be generally positively skewed, the only exception being the negative skewness in the low-strike group for the B&S delta classification. In addition to being mostly positive, skewness across the strike groups is found to be increasing across the group number, such that deep ITM calls exhibit the lowest skewness while deep OTM ones exhibit the highest skewness. With respect to the volatility of option returns, calls in the lowest-strike category tend to earn returns that exhibit less variability, based on the standard error of the distribution, while returns of higher-strike calls are more volatile. Finally, the average number of traded contracts per group appears to be increasing in the group number for groups 1 to 3, since the most ITM calls have the fewest traded contracts and group 3 has the highest number of contracts per option. This monotonic relationship, though, does not hold for the most OTM calls, which appear to have less traded contracts than the options in group 3. These results are not surprising as it's usually the case that options that are the closest to being at-the-money (groups 2 and 3) tend to be more heavily traded and, thus, more liquid than those that are further in or out of the money (groups 1 and 4).

Overall, call returns in the Greek options market are substantially higher than the returns of calls in developed markets. Conforming to theoretical predictions, uncovered positions in calls have earned returns in excess of the underlying asset and increasing in the strike price.

5.5.2 Put Options

The initial put options dataset consisted of 9,212 put observations for the time-period running from January 2004 to January 2007. With p_t denoting the closing price at day t of a put option with strike K and time-to-maturity T, the daily arithmetic return R_p of the option is calculated as:

$$R_p = \frac{p_{t+1} - p_t}{p_t} \tag{5.12}$$

Whenever a specific put does not have traded contracts for two consecutive days or the put price remains the same over this window, the corresponding put return is treated as a missing observation. This results in a reduced dataset of 6,482 put returns.

Similarly to the methodology used in the previous subsection for call options, puts are assigned into four strike groups, using the same three proxies for the options' moneyness. Panel B of Table 5.3 presents the cutoff points, with strike price increasing as the group number increases. This means that group 1 includes put options with the lowest strikes while group 4 includes puts with the highest strikes. However, contrary to calls in Panel A, moneyness for put options in Panel B moves in the opposite direction, with group 1 representing deep OTM puts and group 4 representing deep ITM ones.

Table 5.5 reports the mean, standard error and skewness of put returns across the four strike groups, for all three definitions of moneyness. T-statistics of the average return being different from zero (in brackets) are also reported, as well as the average put beta and average number of traded contracts. The B&S beta of a put option is estimated using the following equation:

$$\beta_p = -\frac{S}{c} N[-\frac{\ln(\frac{S}{X}) + (r - q + \frac{\sigma^2}{2})t}{\sigma\sqrt{t}}]\beta_S \tag{5.13}$$

As can be seen from Panel A in Table 5.5, daily arithmetic put returns in the Greek market have been highly negative and statistically significant for all strike groups, ranging from a minimum of -5.05% for low-strike, deep OTM puts, to a maximum of -2.49% for high-strike, deep ITM ones. Not surprisingly, puts with higher betas (in absolute terms) tend to earn more negative returns than their lower beta counterparts. This makes intuitive sense since puts with high (absolute) betas have relatively low strike prices, representing more levered positions in the underlying asset, and are, thus, perceived as more risky investments. When puts are assigned into groups using the other two moneyness proxies in Panels B and C, results are similar.

Table 5.5 appears to confirm the theoretical prediction of put returns increasing in strike price space. As options move from the low-strike puts in group 1 to the high-strike ones in group 4, average returns increase by becoming less negative, and this monotonic relationship holds for all three option classifications. Moreover, the volatility and skewness of returns is monotonically decreasing in strike price, with skewness remaining positive in all categories. For instance, deep OTM puts in the lowest strike group 1 exhibit the highest standard error (0.42) and skewness (1.81), while high-strike, deep ITM puts in group 4 have returns that are much less volatile (st. error 0.20) and skewed

(0.82). Finally, unlike calls which are more liquid they closest they get to being ATM, put options appear to be more heavily traded when they are OTM. Puts in group 2 are the most liquid, in terms of average traded contracts, with group 1 being the second most liquid category. Deep ITM puts are much less liquid, with average trading volume being 3 or 4 times lower than that of the first two strike groups of OTM puts.

Table 5.5

Summary Statistics for Put Options and their Daily Returns (closing price)

Strike Group	1	2	3	4
		Panel A: K/S		
Put Beta	-36.11	-28.00	-23.04	-16.51
Put Volume (Contract)	115.68	141.99	75.85	36.16
Average Put Return	-0.0505	-0.0482	-0.0363	-0.0249
t-stat	(-6.91)	(-5.58)	(-4.53)	(-3.91)
Put Return St. Dev	0.4198	0.2988	0.2522	0.2008
Put Return Skewness	1.8057	1.1375	0.9235	0.8159
No. of Obs	3,303	1,198	989	991
		Panel B: $\ln(K/S)/\sigma$		
Put Beta	-38.30	-30.50	-21.15	-13.16
Put Volume (Contract)	90.20	148.52	62.22	26.14
Average Put Return	-0.0532	-0.0473	-0.0337	-0.0160
t-stat	(-5.39)	(-6.82)	(-5.72)	(-1.72)
Put Return St. Dev	0.4405	0.3471	0.2379	0.1717
Put Return Skewness	1.8081	1.6253	0.9160	0.5660
No. of Obs	1,994	2,507	1,638	342
		Panel C: B&S delta		
Put Beta	-43.79	-27.67	-21.20	-16.74
Put Volume (Contract)	113.78	132.66	74.19	29.97
Average Put Return	-0.0537	-0.0430	-0.0374	-0.0361
t-stat	(-4.90)	(-6.39)	(-6.18)	(-4.56)
Put Return St. Dev	0.4749	0.3248	0.2445	0.2003
Put Return Skewness	1.7023	1.6737	0.8470	0.6410
No. of Obs	1,875	2,331	1,635	640

This Table tabulates summary statistics of put returns across the period January 2004 to January 2007. Puts have been assigned to moneyness groups according to the cutoff points in Panel B of Table 5.3.

Overall, put options in the Greek market earn negative returns which are decreasing (in absolute terms) as strike price increases, in line with theoretical predictions as well as with empirical findings from other options markets. The returns, however, of short positions in Greek puts are significantly larger than those documented in developed markets, ranging from 400% to 1,342% per year, depending on the strike price and on the moneyness criterion. This implies an asymmetric relationship between returns of calls and puts of similar moneyness, since average put loses significantly outweigh average gains from their corresponding calls.

5.5.3 Option Returns Using the Last Trade Price

The previous two subsections examine returns using closing option prices. In addition to this being a common methodology in the literature, the Athens Derivatives Exchange also states that closing prices are quoted to reflect a representative estimate of the 'true' value of an option contract at the end of each trading day. Although closing levels are typically considered to be a more appropriate proxy for the underlying 'true' value of the option, summary statistics of call and put returns using the price of the last executed trade of the day are also reported.

Table 5.6 presents the results for option returns using last trade prices, grouped under the delta moneyness classification. As can be seen from Panel A, call returns are on average higher for all strike groups when last trade prices are being used, implying that last trade prices are on average lower than closing prices across all groups. For instance, deep OTM calls are found to earn 4.59% per day under the last trade definition, compared to 3.34% under the closing price one. Also, with the exception of deep OTM calls, the standard deviation of returns as well as the B&S betas are also higher under the new proxy. Similarly to call options, puts are found to earn higher (more negative) returns using last trade prices (see Panel B). Only deep ITM contracts in group 4 earn slightly lower returns under the last trade proxy compared to closing prices (-3.55% and 3.61%, respectively).

Table 5.6

Summary Statistics for Options and their Daily Returns (last trade price)

Strike Group	1	2	3	4
	Panel A: Call Options			
Call Beta	19.91	26.76	35.18	38.41
Call Volume (Contract)	7.87	7.45	7.83	14.44
Average Call Return	0.0185	0.0223	0.0366	0.0459
t-stat	(2.76)	(2.26)	(3.13)	(3.53)
Call Return St. Dev	0.21	0.30	0.39	0.45
Call Return Skewness	-0.08	0.91	1.01	1.05
No. of Obs	960	950	1,124	1,210
	Panel B: Put Options			
Strike Group	1	2	3	4
Put Beta	-37.47	-34.55	-25.75	-18.20
Put Volume (Contract)	12.16	8.04	7.04	7.12
Average Put Return	-0.0549	-0.0490	-0.0438	-0.0355
t-stat	(-4.71)	(-4.20)	(-4.43)	(-3.37)
Put Return St. Dev	0.41	0.39	0.29	0.25
Put Return Skewness	1.09	1.92	1.70	0.76
No. of Obs	1,244	1,098	833	565

This Table tabulates summary statistics of the returns of calls and puts in Panels A and B, respectively. The sample period runs from January 2004 to January 2007. Options have been assigned to moneyness groups according to the delta criterion.

Furthermore, the theoretical prediction of option returns increasing as strike price increases is still supported by the results of this alternative methodology, for call as well as for put options. Finally, differences between option returns across different strike groups are larger than those previously reported.

5.6 Estimation of Options' Risk-Adjusted Returns

5.6.1 Single-Factor CAPM

Option returns in the Greek market were found to be surprisingly high between January 2004 and January 2007. Call returns have ranged from around 300% for deep OTM calls to 950% for deep ITM ones, on an annual basis, while short positions in deep ITM and deep OTM puts have earned annual returns of roughly 1,250% and 650%, respectively.

Since options are risky financial assets, standard asset pricing theory predicts that they should earn returns that are commensurate with their systematic risk. The Capital Asset Pricing Model, in particular, expresses the expected excess return of an option as a linear function of the option's beta and the expected market risk-premium:

$$E[R_i - R_f] = \beta_i E[R_{ind} - R_f] \tag{5.14}$$

where R_i is the return of the ith option, β_i is the option's beta, R_{ind} is the return of the FTSE/ASE-20, R_f is the risk-free rate, and $E[\cdot]$ is an expectation operator. Within the CAPM framework, calls that exhibit a higher covariance with the index, as measured by the call's B&S beta, are expected to earn on average higher returns than their lower covariance counterparts. Put options, on the other hand, were shown to have negative betas, which are decreasing in absolute terms as strike price increases. Therefore, puts with more negative betas are expected to earn more negative returns than their lower (in absolute magnitude) beta counterparts.

In order to test the above theoretical predictions, equation (5.15) is regressed separately for calls and for puts across different strike groups

$$R_i - R_f = \eta_0 + \eta_1 \beta_i (R_{ind} - R_f) + \varepsilon \tag{5.15}$$

where η_0 is the intercept term, η_1 is the risk-premium earned by the ith option, and ε is a random error term. Under the CAPM's null hypothesis for this test, the intercept should be statistically indistinguishable from zero (η_0=0) and the risk-premium should be equal to unity (η_1=1), for both calls and puts.

Table 5.7 reports the regression results for all strike groups of calls and puts, under the B&S delta classification. As can be seen from Panel A, the risk-premium of calls, measured as the slope coefficient of (5.15), is significantly positive in all cases, ranging from 0.86 to 1.03. More importantly, risk-premia are found to be statistically indistinguishable from one for ITM calls in groups 1 and 2, and very close to (albeit statistically different from) the theoretical value of unity for OTM calls in groups 3 and 4.

In addition to estimated risk-premia lying close to unity, mostly insignificant intercepts provide further evidence of the CAPM's ability to explain observed call returns. For instance, η_0 is found to be statistically insignificant for ITM as well as for deep OTM calls (groups 1, 2 and 4), with only OTM contracts in group 3 having a significant η_0. Finally, it should be noted that the explanatory power of the model, measured by the Adjusted R^2, is relatively high, ranging from a minimum of 65% (deep OTM) to a maximum of 90% (deep ITM), indicating that the combination of call betas and the market risk-premium can explain a relatively high proportion of the variance of call returns.

Table 5.7

Estimated Regression Coefficients of the CAPM

$$R_i - R_f = \eta_0 + \eta_1\beta_i[R_{ind}-R_f] + \varepsilon$$

Strike Group	1	2	3	4	All
			Panel A: Calls		
η_0	-0.0038	-0.0022	-0.0068	0.0112	-0.0026
t-stat (η_0=0)	(-1.65)	(-1.08)	(-2.39)	(0.86)	(-1.28)
η_1	1.0254	0.9858	0.9476	0.8595	0.9291
t-stat (η_1=0)	(70.22)	(112.55)	(110.20)	(35.91)	(144.91)
t-stat (η_1=1)	(1.74)	(-1.62)	(-6.09)	(-5.87)	(-11.06)
Adj. R^2	0.90	0.85	0.78	0.65	0.75
			Panel B: Puts		
η_0	-0.0285	-0.0143	-0.0115	-0.0118	-0.0184
t-stat (η_0=0)	(-4.90)	(-4.81)	(-4.89)	(-4.54)	(-8.70)
η_1	0.8627	0.9806	0.9931	0.9919	0.9158
t-stat (η_1=0)	(69.37)	(98.48)	(96.04)	(73.43)	(145.16)
t-stat (η_1=1)	(-11.04)	(-1.95)	(-0.67)	(-0.60)	(-13.35)
Adj. R^2	0.72	0.81	0.85	0.89	0.76

This Table tabulates the results from estimating the standard CAPM regression on the daily returns of options written on the FTSE/ASE-20 index. The sample period runs from January 2004 to January 2007. Results for calls and for puts are presented in Panels A and B, respectively.

Panel B of Table 5.7 reports regressions results across all put sub-samples. Risk-premia are statistically indistinguishable from unity for ITM puts (groups 3 and 4) and only marginally different from one in the case of OTM puts in group 2. The η_1 coefficient for deep OTM puts in group 1 is the only exception, since it is found to be significantly lower than its theoretical value of one.

Regarding the intercept terms, η_0 is significantly negative in all cases. Proxied by the regression's intercept, put risk-adjusted returns are monotonically increasing (decreasing in absolute terms) across strikes, with deep OTM puts losing 2.85% and deep ITM ones losing 1.18% on a daily basis. Finally, the Adjusted R^2 is again relatively high, exceeding 89% for deep ITM contracts, indicating that the CAPM's beta has significant explanatory power over put returns.

The common empirical finding of puts earning significantly negative returns after accounting for their risk exposure has frequently been referred to as the 'overpriced puts puzzle'. On the one hand, due to their 'insurance-like' characteristics, it is hardly surprising that index puts are traded at negative risk-premiums since they are negatively correlated with the market. Investors are, therefore, willing to pay more for puts given that they offer payoffs during 'down' markets, and the leverage inherent in option contracts suggests that this risk-premium's magnitude is likely to be large.

On the other hand, most models have so far failed to describe what this risk-premium should be. For instance, Bondarenko (2003) argues that the substantially high returns of selling S&P 500 puts are not only incompatible with two standard asset pricing models, namely the CAPM and Rubinstein's (1976) model, but also that '...*no* model within a fairly broad class of models can possibly explain the put anomaly'. Similar conclusions have been reached by Bollen and Whaley (2004), Buraschi and Jackwerth (2001), and Coval and Shumway (2001).

This anomaly could potentially be explained through the dynamics of supply and demand for calls and puts in an imperfect market. More specifically, Amin, Coval and Seyhun (2004) suggest that portfolio insurance considerations as well as market momentum extracted from past returns have a significant effect on the supply and demand for different types of options. Furthermore, increases in index volatility are likely to drive investors to seek a reduced exposure to the equity market and to bid up the prices of puts relative to calls, while the reverse could be the case if index volatility decreases.

Overall, options in the Greek market appear to be positively related with B&S betas, after controlling for the market risk-premium. In the majority of cases, the slope coefficients of the CAPM regressions are statistically indistinguishable from the theoretical value of unity for both call and put options, while intercepts are equal to zero for calls but significantly negative for puts. The above results seem to imply that the linear risk-return relationship of the standard, single-factor CAPM goes some way into explaining observed option returns in the Greek market.

5.6.2 Extended CAPM

In the mean-variance world of the CAPM, investors are compensated only for bearing the systematic risk stemming from asset returns' covariance with market returns. Recent options literature, though, documents an additional risk-factor being priced in the options market, namely changes in the underlying's volatility. Since options are more valuable when volatility is high, volatility changes should be directly related to option prices and, therefore, option returns. In order to account for this additional risk-factor, an extended version of the CAPM, described in (5.16) is tested:

$$R_i - R_f = \eta_0 + \eta_1 \beta_i (R_{ind} - R_f) + \eta_2 vega_i \Delta\sigma_{imp} \frac{1}{Q_i} + \varepsilon \tag{5.16}$$

where $vega_i$ is the i[th] option's B&S vega, $\Delta\sigma_{imp}$ is the daily change in the FTSE/ASE-20 implied volatility, Q_i is the market price of option i, and η_2 is the corresponding risk-premium. The option's vega is defined as the first derivative of the option's price V with respect to the underlying's volatility σ ($\partial V/\partial\sigma$), and it measures the sensitivity of the option's price to changes in σ. In the context of the above regression specification, σ is defined as the ATM implied volatility[12]. Due to the use of market prices of ATM calls in extracting an estimate of one of the explanatory variables in (5.16), daily returns of ATM calls are excluded from the dependent variable vector $[R_i - R_f]$ since they would provide a near-perfect fit in the regression and would, therefore, introduce some bias into the estimated coefficients.

[12] The ATM implied volatility is estimated by substituting C_{BS} with the actual market price of the nearest-to-the-money call in the Black and Scholes formula, and then solving for the volatility parameter σ.

It should be noted that equation (5.16) represents an extended version of the CAPM in the sense that it requires option returns to be a function of the CAPM market risk-premium as well as of volatility changes. Although changes in the underlying's level and volatility are obviously linked to option returns, the rationale for including these two terms in the regression specification differs significantly. More specifically, the market risk-premium, proxied by the excess return of the index, is assumed to be a determinant of index option returns based on option theory linking the return of the option with that of the underlying, as well as on the rigorous theoretical framework of the CAPM. On the other hand, the risk-premium associated with changes in the volatility of the underlying enters the extended specification based only on option theory and empirical findings suggesting that option prices (and, hence, option returns) are positively related to changes in the underlying's volatility. Therefore, equation (5.16) can only loosely be interpreted as an extended version of the CAPM, since the additional $vega_i \Delta\sigma_{imp}$ term is not the result of a rigorous theoretical derivation. This limitation also suggests that the theoretical volatility risk-premium η_2 is not straightforward to predict, and the subsequent analysis reports only its statistical significance, i.e. statistical difference from zero.

Panels A and B of Table 5.8 report regression results of the extended CAPM for calls and puts, respectively. With respect to call options, it appears that introducing $\Delta\sigma_{imp}$ in the regression does not significantly alter the estimated intercepts η_0 or the slope coefficients η_1. Consistent with theoretical predictions, calls are found to earn a significant volatility risk-premium, with η_2 being significantly positive for all strike groups. Results for put options are not as straightforward at those for calls. Although the slope coefficients η_1 are still very close to unity, intercept terms remain significantly negative. Furthermore, η_2 remains positive for all groups, confirming the theoretical prediction of changes in volatility being positively related to changes in the value of the option. Finally, the correlation between the two main regressors in (5.16), namely between $[R_{ind} - R_f]$ and $\Delta\sigma_{imp}$, is -0.10 over the sample period. Although a negative correlation between the market risk-premium and changes in index volatility is to be expected since it is a well-documented empirical finding that index volatility is systematically higher following a negative index return compared to a positive return of similar magnitude, it could be argued that the low level of dependence between the two explanatory variables suggests that the regression results are relatively free of collinearity concerns.

Table 5.8

Estimated Regression Coefficients of the extended CAPM

$$R_i - R_f = \eta_0 + \eta_1 \beta_i [R_{ind} - R_f] + \eta_2 vega_i \Delta\sigma_{imp} \frac{1}{Q_i} + \varepsilon$$

Strike Group	1	2	3	4	All
		Panel A: Calls			
η_0	-0.0039	-0.0023	-0.0075	0.0038	-0.0038
t-stat (η_0=0)	(-1.77)	(-1.24)	(-2.86)	(0.31)	(-2.00)
η_1	1.0152	0.9979	0.9623	0.9023	0.9496
t-stat (η_1=0)	(72.97)	(121.56)	(120.68)	(40.81)	(159.71)
t-stat (η_1=1)	(1.09)	(-0.26)	(-4.72)	(-4.42)	(-8.48)
η_2	6.6113	0.9552	0.6889	1.0664	0.8197
t-stat (η_2=0)	(7.76)	(17.92)	(24.20)	(12.02)	(34.71)
Adj. R^2	0.91	0.87	0.81	0.71	0.79
corr(R_{ind}-R_f, $\Delta\sigma_{imp}$)	-0.01	-0.10	-0.09	-0.23	-0.11
		Panel B: Puts			
η_0	-0.0298	-0.0160	-0.0120	-0.0112	-0.0194
t-stat (η_0=0)	(-5.20)	(-5.42)	(-5.14)	(-4.34)	(-9.28)
η_1	0.8661	0.9798	0.9920	0.9907	0.9175
t-stat (η_1=0)	(70.57)	(99.53)	(97.17)	(73.71)	(147.32)
t-stat (η_1=1)	(-10.91)	(-2.05)	(-0.78)	(-0.69)	(-13.24)
η_2	0.0865	0.0990	0.2007	0.3056	0.0924
t-stat (η_2=0)	(7.34)	(7.44)	(6.62)	(2.88)	(13.15)
Adj. R^2	0.73	0.81	0.85	0.90	0.77
corr(R_{ind}-R_f, $\Delta\sigma_{imp}$)	0.03	-0.04	-0.01	-0.01	-0.08

This Table tabulates the results from estimating the extended CAPM on the daily returns of options written on the FTSE/ASE-20 index. The estimated coefficients for calls and for puts are presented in Panels A and B, respectively. The sample period runs from January 2004 to January 2007, and options have been assigned to moneyness groups based on the cutoff points of Table 5.3. The last row of each Panel tabulates the correlation between the two dependant variables, namely between the excess return of the market and the daily change in implied volatility.

5.7 Straddles

In the previous Sections, options in the Greek market were found to earn significantly higher returns than those traditionally earned by options in developed markets. Daily returns of uncovered, long positions in calls have been positive, increasing across strike price, and significantly higher than the relatively large market risk-premium for the sample period. Uncovered, long positions in puts, on the other hand, have earned significantly negative returns, which are also increasing (decreasing in magnitude) as strike price increases. Although a positive relationship was documented between the above returns and the CAPM's systematic variance, as well as changes in the underlying's volatility, significantly negative intercepts for put options suggest that additional risk-factors might be priced in the Greek market.

This Section shifts the attention from individual options to option portfolios. More specifically, the Section examines returns of portfolios formed by combining long/short positions in calls and puts, in a way that ensures the resulting portfolios have zero exposure to at least one of the widely accepted sources of risk in the options market.

5.7.1 Delta-Neutral Straddles

Delta-neutral portfolios are formed by combining long positions in calls and puts of the same moneyness, with moneyness defined as $(1 - K_c/e^{rT}S)$ and $(K_p/e^{rT}S - 1)$ for calls and for puts, respectively. In order to estimate the weights w_c and w_p of the portfolio's value that correspond to investing in calls and puts, respectively, the following two equations are simultaneously solved. The first equation stems from the straddle's delta ($delta_s$), which is a linear combination of individual options' deltas ($delta_c$ and $delta_p$), being equal to zero, while the second one reflects the fact that the possible combinations are restricted to only long positions in same-moneyness options. Obviously, since B&S deltas are by definition positive for call options and negative for put options, there is only one possible combination that satisfies both of these conditions.

$$delta_s = w_c delta_c + w_p delta_p = 0$$

$$w_c + w_p = 1$$

Since both beta and delta are measures of an option's sensitivity to changes in the value of the market index, they are effectively proxies for the same type of market risk. In other words, delta-neutral portfolios of index options can also be considered as beta-neutral positions. In the CAPM world, systematic variance, proxied by a security's beta, is the only source of risk that is priced in the market. Therefore, a delta-neutral index straddle has zero exposure to systematic risk, and should earn the risk-free rate of return.

Table 5.9 presents descriptive statistics for the daily returns of delta-neutral straddles across the four moneyness groups. These groups are formed such that moneyness increases as options move from the first group to the fourth one, with group 1 including long positions in deep OTM options and group 4 including deep ITM ones. As can be seen from the Table, delta-neutral portfolios are found to earn relatively low returns which are increasing across moneyness. More specifically, deep OTM straddles lose around 3 basis points, while deep ITM portfolios earn 38 basis points on a daily basis, with average returns increasing as we move from group 1 to group 4. More importantly, though, straddle returns are statistically indistinguishable from the risk-free rate for all moneyness groups, confirming theoretical predictions that option combinations that are immune to changes in the value of the underlying should be considered as risk-free and, therefore, have returns equal to the daily risk-free rate. In addition, straddle returns exhibit low volatility (roughly 1% across all groups), slightly negative skewness and negative excess kurtosis.

Overall, results from examining delta-neutral straddles provide some support for the validity of the Black-Scholes model as well as the CAPM is the Greek options market. However, although the theoretical prediction that portfolios with zero delta-risk should offer returns that are equal to the risk-free rate is empirically confirmed, it should be noted that delta-neutral straddles are potentially exposed to other sources of risk. For instance, long positions in these straddles typically have high volatility betas and are, therefore, more profitable in periods of high volatility.

Table 5.9

Daily Returns of Delta Neutral Straddles

Moneyness Group	m < -0.03	-0.03 < m < 0	0 < m < 0.03	0.03 < m
Average	-0.0003	0.0013	0.0017	0.0038
St. Error	0.01	0.01	0.01	0.01
t-stat	(-0.03)	(0.13)	(0.18)	(0.44)
Median	0.0002	0.0014	0.0019	0.0040
Skewness	-0.17	-0.22	-0.18	-0.12
Kurtosis	1.22	1.11	1.09	0.14
No of Obs.	2,347	1,852	1,341	1,022

This Table presents summary statistics of the daily returns of delta-neutral straddles across the sample period from January 2004 to January 2007. The t-statistics in brackets refer to the average daily straddle return being statistically significant from the daily risk-free rate. The daily risk-free rate (proxied by Euribor) was roughly equal to 1 basis point (or equivalently 2.72% per annum) during the sample period.

5.7.2 Delta and Vega Neutral Straddles

In the previous subsection, combinations of long positions in calls and puts were formed such that the exposure of the overall position to market risk, given as the weighted average of the options' deltas, was zero. Although this methodology ensures that the resulting portfolio is effectively immune to changes in the value of the FTSE/ASE-20, such call-put combinations are not necessarily risk-free. An additional, widely recognised source of risk in the options market refers to changes in the level of the underlying's volatility until the option's expiration, and it is measured by the option's vega ($\partial V/\partial \sigma$). Intuitively, since the value of an option is positively related to the future volatility σ of the underlying, changes in volatility will affect the option's price and, consequently its expected return.

The methodology of Liu (2007) is followed to create delta and vega neutral portfolios by combining long positions in the underlying and in puts with short positions in calls of similar moneyness, with moneyness defined as $(1 - K_c/e^{rT}S)$ and $(K_p/e^{rT}S - 1)$ for calls and puts, respectively. More specifically, on each calendar day, delta and vega neutral straddles are formed by buying one unit of the index and w_c units of the call, while selling w_p units of the put, the moneyness of which is the closest to the call's moneyness. In order for the straddle's exposure to delta and vega risk to be zero, the following conditions must be met:

$$delta_s = 1 + w_c delta_c + w_p delta_p = 0$$

$$vega_s = w_c vega_c + w_p vega_p = 0$$

Obviously, the delta of the underlying is equal to one and its vega is zero. Also, calls have positive deltas and puts have negative ones, but both options have positive vegas. Therefore, w_p must be positive and w_c negative to ensure that $delta_s = vega_s = 0$. In total, 6,562 delta and vega neutral straddles are formed following the previously described methodology. The average difference in moneyness between calls and puts is 0.0126, with 69% of straddles including options with moneyness levels that differ by a maximum of 0.01. Straddles are then assigned to four groups based on their moneyness, with group 1 including combinations of options that are deep OTM and group 4 including deep ITM options. Panel A of Table 5.10 presents summary statistics for the daily returns of delta and vega neutral straddles, across the four moneyness groups.

The null hypothesis of interest is that straddles with zero risk-exposure to market movements and to volatility changes must earn the risk-free rate. Indeed, it is found that daily straddle returns for all moneyness groups are statistically indistinguishable from zero, as well from the daily risk-free rate. Average returns are increasing across moneyness, with deep OTM straddles earning negative returns (0.09%) and deep ITM straddles earning the highest positive returns (0.35%).

Although these results are in line with theoretical predictions, it should be noted that this methodology suffers from a relatively well known limitation. More specifically, B&S delta and vega are measures of local sensitivity, referring to expected changes in the option's price for a marginal change in the index's level and volatility, respectively. Therefore, straddles created in the above way will be delta and vega neutral only instantaneously and with respect to very small changes in the FTSE/ASE-20 and its volatility. In order to ensure near-zero risk-exposure, straddles have to be rebalanced regularly, at the obvious expense of higher transaction costs. For instance, Liu (2007) argues that such portfolios '... start off delta and vega neutral, but the neutrality is unlikely to hold in one week's time'. In this study, straddle returns are examined at a daily frequency in an attempt to minimise the impact of changing delta/vega across our holding period and it is found that, even without considering the substantially higher transaction costs, the null hypothesis of straddles earning the risk-free rate cannot be rejected.

Table 5.10

Daily Returns of Delta and Vega Neutral Straddles

Moneyness Group	$m < -0.03$	$-0.03 < m < 0$	$0 < m < 0.03$	$0.03 < m$
Panel A: Summary Statistics for Delta and Vega Neutral Straddles				
Average	-0.0009	0.0006	0.0009	0.0035
St. Error	0.02	0.01	0.01	0.02
t-stat	(-0.06)	(0.04)	(0.06)	(0.22)
Median	-0.0003	0.0010	0.0010	0.0038
Skewness	0.19	0.07	0.40	0.04
Kurtosis	2.85	1.70	2.00	0.60
No of Obs.	2,347	1,852	1,341	1,022
Panel B: Estimated Regression Coefficients $E[R_s] = \eta_0 + \eta_1 E[R_{ind}] + \eta_2 \Delta\sigma_{imp} + \varepsilon$				
η_0	-0.003	-0.002	-0.004	-0.006
t-stat	(-2.17)	(-1.22)	(-1.64)	(-2.03)
η_1	0.050	-0.053	-0.005	-0.021
t-stat	(1.71)	(-1.58)	(-0.11)	(-0.39)
η_2	0.012	0.014	0.023	0.047
t-stat	(1.62)	(1.63)	(2.07)	(3.23)

Panel A of this Table tabulates summary statistics of the returns of delta and vega neutral straddles for the sample period January 2004 to January 2007. Panel B tabulates the results of estimating the extended CAPM on the daily returns of the straddles.

In order to determine whether straddles are indeed delta and vega neutral, portfolio returns are regressed against the future changes in the underlying FTSE/ASE-20 and the future changes in implied volatility that corresponds to the period until the options' expiration, using equation (5.17).

$$E[R_s] = \eta_0 + \eta_1 E[R_{ind}] + \eta_2 \Delta\sigma + \varepsilon \tag{5.17}$$

where R_s is the daily straddle return, and $\Delta\sigma$ is the daily change in B&S implied volatility. Under the assumption of delta and vega neutrality, the null hypothesis is that $\eta_0 = \eta_1 = \eta_2 = 0$, and Panel B of Table 5.10 reports the regression results across the four moneyness groups.

The first thing to notice is that η_1 is insignificant for groups 2 to 4 and borderline insignificant for deep OTM options in group 1. In effect, straddles across all moneyness categories remain approximately delta-neutral during the trading day of interest and are, thus, not affected by changes in the level of the underlying index. However, although delta neutrality seems to hold, not all straddle types are vega neutral. The vega-neutrality coefficients η_2 for OTM groups 1 and 2 are marginally insignificant with t-stats equal to 1.62 and 1.63, respectively, while ITM straddles in groups 3 and 4 have significant η_2s (t-stats are 2.07 and 3.23, respectively).

In summary, delta and vega neutral straddles appear to earn returns that are statistically indistinguishable from the daily risk-free rate, irrespective of their moneyness. This finding supports the theoretical prediction that positions that are immune to the two most common sources of risk in the options market, namely changes in the level of the underlying and changes in the underlying's volatility, should earn the risk-free rate. In addition, although straddle returns have been calculated using closing option prices, it should be mentioned that the results remain unchanged even when last trade prices are used instead. However, when interpreting these results, one should have in mind that, despite the fact that all straddles appear to be delta-neutral in the one-day period, ITM positions are subject to some vega risk.

5.7.3 Risk-Reversals

A significant body of the options literature documents that Risk Neutral Densities which are inferred from option prices exhibit significant negative skewness (see, for instance, Jackwerth (2000)). One explanation that has been proposed for this well documented finding is options mispricing, especially overpricing of OTM puts. If the options market implies a significant negative skewness that is not consistently related to a subsequent fat-tailed realized distribution, trading strategies that short overpriced OTM puts are likely to earn returns that are in excess of their risk.

The methodology of Liu (2007) is followed to construct risk-reversals, which are portfolios that essentially place a bet on the non-realization of option-implied negative skewness, by exploiting the difference in price between call and put options of the same moneyness. On the first trading day of each month (i.e. the Monday after the third Friday of each month), pairs of calls and puts with similar absolute deltas are identified, and a risk-reversal is created by shorting one contract of the more expensive option while buying one contract of the cheaper one. It should be noted that, in forming these portfolios, call-put combinations are considered when their absolute deltas differ by a maximum of 0.04. Buying the call and shorting the put, then, is a long risk-reversal while the reverse combination is characterized as a short risk-reversal.

Contrary to the methodology followed in the previous subsections, this Section's main focus is on holding-to-expiration payoffs rather than on daily returns. Although the initial value of the portfolio is positive by definition, in the sense that the investor receives the price differential between the two option contracts, there are three possible terminal payoff types, depending on the value of the underlying index on expiration relative to the strikes range $[K_p, K_c]$. For instance, in the case of a long risk-reversal with OTM options, if the price S_T of the FTSE/ASE-20 on expiration remains in the range bounded by the strikes, the contracts have equal terminal payoffs of the opposite sign (i.e. the position's net payoff is zero), and the investor profits from keeping the options' premia differential. If S_T falls below the lower put strike K_p, the call expires worthless while there is a negative payoff from the short position in the put. Finally, if S_T rises above the higher call strike K_c, there is a positive payoff from the long position in the call while the put will not be exercised. .

Following the above methodology, 189 holding-to-expiration risk-reversals are formed across different moneyness levels by combining calls and puts as long as the absolute difference of their deltas is less than 4%. Since puts are usually more expensive than same-moneyness calls, 129 out of the total of 189 portfolios are long positions, while the remaining 60 are short risk-reversals. More than half of the long positions

(58.14%) have zero payoffs, leaving the investor profiting from the difference in the options' premia that was received when the portfolio was formed. In addition, 32.56% of long risk-reversals achieve positive payoffs on expiration, and only 9.30% have negative terminal payoffs. Regarding short risk-reversals, 31.67% have zero payoffs upon expiration, while most of them (46.67%) achieve positive payoffs. Finally, only 21.67% of the 60 short portfolios have negative payoffs, indicating that the level S_T of the index at maturity was above the range of strikes $[K_p, K_c]$ for a small proportion of our sample.

Table 5.11

Risk-Reversals

Delta Group	delta < 0.25	0.25 < delta < 0.50	0.50 < delta < 0.75	0.75 < delta
	Panel A: Summary Statistics for Daily Returns of Risk-Reversals			
Average	0.03	0.12	0.36	0.02
t-stat	(0.02)	(0.04)	(0.04)	(0.06)
Median	0.21	0.15	0.15	0.03
No of Obs.	386	1,701	1,059	142

Panel B: Returns for Holding-to-Maturity Risk-Reversals

	All	Long	Short
	189 (100%)	129 (68.25%)	60 (31.75%)
Zero Payoffs	94 (49.74%)	75 (58.14%)	19 (31.67%)
Positive Payoffs	70 (22.22%)	42 (32.56%)	28 (46.67%)
Negative Payoffs	25 (13.23%)	12 (9.30%)	13 (21.67%)

Panel A of this Table tabulates summary statistics of the daily returns of risk-reversals across four moneyness groups. The sample runs from January 2004 to January 2007. Panel B reports the returns of holding-to-maturity positions.

In summary, risk-reversals are trading strategies designed to profit from betting against the options market's expectation of a significant decline in the value of the underlying index. For the sample period, investing in these portfolios in the Greek market appears to offer significant real profit opportunities, without adjusting for risk, since the majority of examined positions achieve positive cashflows, stemming from both the options' premia differential as well as from payoffs upon expiration. However, it has to be noted that the period from January 2004 to January 2007 is a relatively short one, and it corresponds to a bullish Greek market, with the large capitalization FTSE/ASE-20 appreciating by more than 28% annually. Consequently, the above mentioned results should be interpreted cautiously, since they might be attributed to a 'small sample problem', in the sense that the market priced options to reflect the true probability of a market fall which never materialized in the somewhat short period under examination. The possibility of a small sample biasing these results does not necessarily imply that Greek OTM puts are correctly priced, but highlights the fact that, should the market had followed a more negatively skewed distribution, terminal payoffs from investing in risk-reversals would have been much lower or potentially negative overall.

5.8 Comparison with Empirical Findings from Developed Markets

In order to put the above findings into context, option returns in the Greek market are compared to those observed in developed markets. Although the high magnitude of returns of individual calls and puts written on the FTSE/ASE-20 seems puzzling at a first glance, compared to option returns written on the S&P 500 or the FTSE100, it is concluded that they are not necessarily inconsistent with traditional option pricing models. Furthermore, returns of various trading strategies, such as delta and vega neutral straddles, indicate that risk-return theoretical predictions are strongly supported in the Greek market, similarly to its UK and US counterparts.

As has already been mentioned in Section 5.1.2, the efficiency of the emerging ADEX is evaluated by comparing the results of this Chapter with stylized facts from traditional, developed options markets. Previous studies examining the US, in particular, have shown that calls do not usually earn significant risk-adjusted returns (proxied by CAPM intercepts) while risk-adjusted put returns are typically found to be significantly negative (see Bondarenko (2003), Broadie, Chernov and Johannes (2009), Coval and Shumway (2001), and Santa-Clara and Saretto (2009) for an analysis of the CAPM on US options). In addition, most empirical papers suggest that risk-immune straddles in the US only earn the risk-free rate, although some deviations have been documented depending on the straddle's moneyness and the period examined (see Bakshi and Kapadia (2003), Broadie, Chernov and Johannes (2009), and Coval and Shumway (2001) for an examination of delta-neutral position in the US). This finding supports market efficiency in the sense that an investor establishing a zero-risk position is compensated by earning only the risk-free rate.

5.8.1 Naked Options Positions

First, call options in developed markets have been found to earn relatively high average returns. For instance, Coval and Shumway (2001) focus on options written on the S&P 500 index from January 1986 to October 1995, and report weekly call returns ranging from 1.48% for deep ITM calls to 5.13% for deep OTM ones. Supporting theoretical predictions, these returns are in excess of the underlying's rate of appreciation for the same time period and monotonically increasing as strike price increases. On the other hand, Driessen and Maenhout (2006) examine returns of options written on the FTSE100 from April 1992 to June 2001, and find that, in contrast to S&P options, returns of UK calls have been significantly lower. Average weekly returns of short-term FTSE100 calls range from 0.28% for ATM options to 0.04% for deep OTM ones, while the theoretical monotonic relationship between returns and moneyness is not supported. As was discussed in Section 5.5.1, returns of calls written on the FTSE/ASE-20 have been significantly higher than those previously documented in the US and the UK markets. More specifically, deep ITM calls earn 5.75% per week, while deep OTM ones earn around 19.05% per week, with call returns strictly increasing across strikes. As has already been noted, though, the fact that average returns of Greek call options are four times higher compared to the US, and more than fifty times the magnitude of UK call returns of similar moneyness, might be at least partially explained by the rapid growth of the underlying FTSE/ASE-20 during the 2004-2007 sample period.

A well documented finding in the related literature refers to the fact that put options tend to earn on average higher returns (in absolute terms) than calls of similar moneyness. This asymmetry is highlighted by Coval and Shumway (2001) for the US market, with puts written on the S& 500 losing between -14.56% for deep OTM puts to -6.16% for deep ITM ones. Confirming theoretical predictions, these put returns are below the risk-free rate, as well as increasing (becoming less negative) as strike price increases (see also Bondarenko (2003) and Broadie, Chernov and Johannes (2009) for returns of S&P 500 puts within different sample periods). With respect to the UK, Driessen and Maenhout (2006) find that puts written on the FTSE100 have highly negative returns, ranging from -6.86% for short-term, deep OTM options to -4.58% per week for deep ITM ones. In contrast to FTSE100 calls, puts support the theoretical prediction of strictly increasing returns (decreasing in magnitude) across moneyness, while the difference between put returns in the US and in the UK is significantly smaller than the one documented for calls. Section 5.5.2 of this Chapter reports that puts written on the FTSE/ASE-20 lose between -25.25% (deep OTM) and -12.45% (deep ITM) on a weekly basis, while, similarly to results for developed markets, put returns in Greece become less negative as strike price increases. Overall, writing put options on the FTSE/ASE-20 results in higher average returns compared to same-moneyness S&P 500 or FTSE100 puts, with put returns in the Greek market being closer to results from the US rather than the UK options market.

The fact that options are found to consistently earn very high returns, with the most extreme case typically being returns to writing deep OTM puts, has led some authors to describe options returns as 'puzzling'. However, Broadie, Chernov and Johannes (2009) argue that, unless they are compared to a reasonable benchmark, it is difficult to conclude whether high option returns constitute in fact a paradox. The significance of observed returns can be examined by focusing on risk-adjusted estimates, proxied by the intercepts of CAPM regressions on option returns. Under standard CAPM theory, alphas are expected to be statistically indistinguishable from zero, since expected returns are only compensating for bearing systematic risk. Focusing on the Greek market, this theoretical prediction is supported in the case of calls which have insignificant alphas. Put returns appear to be relatively puzzling, since after controlling for their systematic risk, intercepts remain statistically significant across all moneyness categories. These results are similar to Broadie, Chernov and Johannes' (2009) examination of risk-adjusted returns of puts written on the S&P500, who report statistically significant CAPM alphas, ranging from -51.72% for deep OTM puts to -24.60 for deep ITM ones (on a monthly basis). Also note that, these results remain unchanged, even after incorporating changes in the underlying's volatility as an additional explanatory factor in the extended CAPM, indicating that additional factors are potentially priced in the Greek market, a conclusion that is consistent with the explanation proposed by Broadie, Chernov and Johannes (2009) for the US market.

5.8.2 Option Strategies

After examining individual option returns, the focus moves to returns of various option portfolios. First, a trading strategy that has received a fair amount of attention in the related literature explores returns of delta-neutral combinations of options which, under the CAPM's assumptions, should be equal to the risk-free rate. The intuition behind this methodology is that, since these straddles are formed such that they are essentially zero-delta (or, equivalently, zero-beta), they have no exposure to risk from market movements and, consequently, should earn the risk-free rate. However, Coval and Shumway (2001) as well as Broadie, Chernov and Johannes (2009) report that delta-neutral straddles which are formed by combining calls and puts written on the S&P 500 have, in fact, statistically significant returns in their respective sample periods. For instance, Coval and Shumway (2001) find that ATM straddles lose around 3% on a weekly basis, with straddle returns generally increasing (becoming less negative) as strike price increases. In contrast, zero-delta straddles in Greece are found to have insignificant returns, irrespective of their moneyness, indicating that the CAPM's market-risk alone goes some way into explaining the return characteristics of options combinations.

It has been suggested that the significant returns of delta-neutral straddles in developed markets are due to the fact that, although these straddles are theoretically immune to changes in the value of the underlying, they might still be exposed to other sources of risk. The attention that volatility risk has received in recent studies prompts Liu (2007) to focus on delta and vega neutral straddles, combining long positions in the underlying and a put with a short position in a call of similar moneyness. Since these straddles have zero exposure to the two most commonly accepted sources of risk in the options market, namely changes in the value of the underlying as well as changes in the underlying's volatility, they are expected to earn the risk-free rate. However, Liu (2007) examines a sample of options written on the FTSE100 from January 1996 to April 2000, and finds that, while weekly returns of delta and vega neutral straddles are insignificant for ATM and ITM combinations, OTM and deep OTM straddles have significantly negative returns. Moreover, Liu argues that one potential explanation for the above mentioned results might be that, since delta and vega are estimated as local sensitivities, the straddles' neutrality is unlikely to hold across the return-estimation period of one week. As was discussed in Section 5.7.3 of this Chapter, delta and vega neutral straddles on the FTSE/ASE-20 have returns that are slightly negative and increasing across strikes. More importantly, though, straddle returns across all moneyness levels are statistically indistinguishable from the risk-free rate, supporting theoretical predictions.

Finally, Liu (2007) examines whether the common finding that deep OTM puts are typically more expensive than calls of similar moneyness, presents any real profit opportunities in the UK market. She estimates holding-to-expiration returns of FTSE100 risk-reversals, which are option combinations designed to profit from a bet against the negative market skewness implied by deep OTM option prices, and finds that, without accounting for risk, these trading strategies have been significantly profitable during her sample period. More specifically, OTM puts are found to be more expensive than similar moneyness calls in 81% of the positions examined, with more than 80% of these positions resulting in positive returns, stemming from the options' premia differential as well as the positive terminal net payoff of the portfolio. Section 5.7.4 reports similar results for risk-reversals on the FTSE/ASE-20, with 85%, on average, of all portfolios having positive payoffs. However, the possibility of a small sample problem biasing the results, especially in the bullish Greek market, has to be kept in mind when interpreting the above findings.

5.9 Conclusion

This Chapter has examined the efficiency of the emerging market of the Athens Derivatives Exchange compared to developed options markets, from the perspective of returns that are commensurate with underlying risks. It is shown that returns of individual options in Greece are not inconsistent with empirical findings from developed options markets, such as the US and the UK. In addition, returns of delta and delta/vega neutral straddles are found to be statistically indistinguishable from the risk-free rate, implying that returns of these trading strategies are in line with theoretical predictions, with p-values even higher than those documented in traditional, developed markets.

These results appear to reject the hypothesis of ADEX exhibiting a lower level of efficiency, attributed to the relatively high transaction costs and illiquid trading in the Greek options market, compared to the US. Santa-Clara and Saretto (2009) document a potential mispricing in S&P options that results in various option strategies earning abnormally high returns relative to their risk. However, these profit opportunities are allowed to persist instead of being arbitraged away due to the relatively high bid-ask spread, as well as the strict margin requirements in the US market. Following this line of thought, one might expect that the Greek market would offer a greater scope for options mispricing, since exploiting these profit opportunities would be even more costly for a typical investor due to the significantly higher bid-ask spreads as well as to thinner trading.

In order to put the significant difference in trading costs between developed and emerging markets into context, one should consider that trading volume in Greece is dramatically lower than, for instance, the US. During 2006, slightly less than 600,000 FTSE/ASE-20 option contracts were traded in ADEX, while the respective volume for S&P500 options traded in CBOE exceeded 104 millions. In addition, the fees charged by the exchange for each transaction are higher in the Greek market, with ADEX charging market-makers around €0.20 per trade (depending on the option's moneyness) while CBOE charging around $0.20 per trade (depending on total contracts traded). [13]

However, the relative efficiency of ADEX cannot be rejected, since trading strategies do not appear to offer significant profit opportunities in this emerging options market, even without accounting for transaction costs. This seems to indicate that the Greek market exhibits a degree of efficiency comparable to that of developed markets, at least with respect to the absence of opportunities for abnormal profits in excess of the underlying risks.

[13] Chicago Board Options Exchange, 2006 Annual Report

Appendix B

Theoretical Expected Option Returns

Given that options represent levered positions in the underlying asset, the Black-Scholes/CAPM framework predicts that options' betas should be greater in absolute terms compared to the underlying's beta. Calls written on a market index (which has a beta of unity and a positive expected risk-premium) will have positive betas higher than one and, therefore, theoretically expected call returns are positive and higher than the underlying index's returns. Accordingly, index puts are negatively related to changes in the underlying's level, i.e. put betas are negative, so they should theoretically have expected returns that are below the risk-free rate as they are essentially instruments for hedging systematic risk.

Although it could be argued that the Black-Scholes as well as the CAPM assumptions are potentially restrictive, Coval and Shumway (2001) demonstrate that these theoretical predictions can be obtained under much weaker assumptions. The following discussion is based on the derivation of Propositions 1 and 2 in the original Coval and Shumway (2001) paper.

The Coval and Shumway general setting only assumes the existence of a stochastic discount factor m that prices all assets through (B.1) and that this stochastic discount factor is negatively correlated with the price of the underlying,

$$E[R_i^G \times m] = 1 \tag{B.1}$$

where R_i^G is the gross return of asset i and m is strictly positive. Then, Proposition 1 describes the expected return of a call option that is written on security i (quote from the original paper):

Proposition 1 *If the stochastic discount factor is negatively correlated with the price of a given security over all ranges of the security price, any call option on that security will have a positive expected return that is increasing in the strike price.*

Proof Assume that the underlying security has a distribution $f(y)$ and that the joint distribution of the security and the stochastic discount factor is given by $f(y,z)$. Then the expected gross return $E[R_c^G]$ of a call option with strike K can be expressed as:

$$E[R_c^G(K)] = \frac{\int_{S=K} (S-K) f(S) \partial S}{\int_{q=0} \int_{S=K} q(S-K) f(S,q) \partial S \partial q} \tag{B.2}$$

where S is the price of the underlying. The expected net return $E[R_c^N]$ of the call option can be written as:

$$E[R_c^N(K)] = E[R_c^G(K)] - 1 = \frac{\int_{S=K} (S-K)(1-E[m \mid S]) f(S) \partial S}{\int_{S=K} (S-K) E[m \mid S] f(S) \partial S} \tag{B.3}$$

Coval and Shumway (2001) then show that the derivative $\frac{\partial E[R_c^N(K)]}{\partial K}$ of expected net call returns with respect to the strike price is:

$$\frac{\partial E[R_c^N(K)]}{\partial K} =$$

$$\frac{\int_{S=K} (S-K) \frac{f(S)}{1-F(K)} \partial S \int_{S=K} E[m \mid S] \frac{f(S)}{1-F(K)} \partial S - \int_{S=K} (S-K) E[m \mid S] \frac{f(S)}{1-F(K)} \partial S}{[\int_{S=K} (S-K) E[m \mid S] \frac{f(S)}{1-F(K)} \partial S]^2} \tag{B.4}$$

where $F(S)$ is the cumulative density function corresponding to $f(S)$. The numerator in (B.4) is the negative of the covariance between $(S - K)$ and m, conditional on S > K, i.e. the option being in the money:

$$-\mathrm{cov}(E[m \mid S], (S-K) \mid S > K] = E[m \mid S > K] \times E[S - K \mid S > K] - E[E[m \mid S](S-K) \mid S > K] \tag{B.5}$$

The assumption that the stochastic discount factor m is negatively correlated with the security price S implies that the expression in (B.5) will be positive. In other words, expected net call returns $E[R_c^N(K)]$ will be increasing in strike price K since their first derivative with respect to K was shown to be positive, $\frac{\partial E[R_c^N(K)]}{\partial K} > 0$. This relation, combined with the fact that a call option with a zero strike price will have an expected return equal to that of the underlying, suggests that expected call returns should theoretically be positive, higher than the underlying's expected return, and monotonically increasing across strikes.

The corresponding Proposition 2 describes expected put returns (quote from the original paper):

Proposition 2 *If a stochastic discount factor is negatively correlated with the price of a given security over all ranges of the security price, any put option on that security will have an expected return below the risk-free rate that is increasing in the strike price.*

Proof Let the expected net return $E[R_p^N(K)]$ of a put option with strike K written on an underlying which has a distribution $f(y)$ and a joint distribution with the stochastic discount factor $f(y,z)$ be given by:

$$E[R_p^N(K)] = \frac{\int_{S=l}^{S=K} (K-S)(1-E[m \mid S])f(S)\partial S}{\int^{S=K} (K-S)E[m \mid S]f(S)\partial S} \tag{B.6}$$

The first derivative $\frac{\partial E[R_p^N(K)]}{\partial K}$ of expected net put returns with respect to the strike price is:

$$\frac{\partial E[R_p^N(K)]}{\partial K} =$$

$$\frac{\int^{S=K}(K-S)E[m \mid S]f(S)\partial S \int^{S=K} f(S)\partial S - \int^{S=K}(K-S)f(S)\partial S \int_{S=K} E[m \mid S]f(S)\partial S}{[\int^{S=K}(K-S)E[m \mid S]f(S)\partial S]^2} \qquad \text{(B.7)}$$

Using the same reasoning as in Proposition 1, the numerator in (B.7) is proportional to the covariance between (K - S) and the stochastic discount factor m, conditional on $S < K$, i.e. the option being in the money.

$$\operatorname{cov}(E[m \mid S],(K-S) \mid S<K] = E[E[m \mid S](S-K) \mid S<K] - E[m \mid S<K] \times E[S-K \mid S<K] \qquad \text{(B.8)}$$

As the strike price approaches infinity, the expected net put return approaches the risk-free rate. Also, the positive first derivative $\frac{\partial E[R_p^N(K)]}{\partial K}$ of put returns with respect to strike price results in net put returns increasing monotonically as K increases. Therefore, expected put returns should theoretically be lower than that of the underlying (although not necessarily negative) and increasing across strike price space.

References

Amin, Kaushik, Joshua Coval, and Nejat Seyhun, 2004, Index Option Prices and Stock Market Momentum, *Journal of Business* 77, 835-873

Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta-Hedged Gains and the Negative Market Volatility Risk Premium, *Review of Financial Studies* 16, 527-566

Black, Fisher, and Myron Scholes, 1973, The pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637-654

Bollen, Nicolas, and Robert Whaley, 2004, Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?, *Journal of Finance* 59, 711-753

Bondarenko, Oleg, 2003, Why are Put Options so Expensive?, *University of Illinois working paper*

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2009, Understanding Index Option Returns, forthcoming in *Review of Financial Studies*

Buraschi, Andrea, and Jens Jackwerth, 2001, The Price of a Smile: Hedging and Spanning in Option Markets, *Review of Financial Studies* 14, 495-527

Christie-David, Rohan, and Mukesh Chaudhry, 2001, Coskewness and Cokurtosis in Futures Markets, *Journal of Empirical Finance* 8, 55-81

Corrado, Charles, and Tie Su, 1996, Skewness and Kurtosis in S&P500 Index Returns Implied by Option Prices, *Journal of Financial Research* 19, 175-192

Coval, Joshua, and Tyler Shumway, 2001, Expected Option Returns, *Journal of Finance*, 56, 983-1009

Dittmar, Robert, 2002, Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns, *Journal of Finance* 57, 368-403

Fama, Eugene, and Kenneth French, 1992, The Cross-Section of Expected Returns, *Journal of Finance* 59, 427-465

Fama, Eugene, and James MacBeth, 1973, Risk, Return and Equilibrium: Empirical Tests, *Journal of Political Economy* 81, 607-636

Fang, Hsing, and Tsong-Yue Lai, 1997, Cokurtosis and Capital Asset Pricing, *Financial Review* 32, 293-307

Floros, Christos, and Dimitrios Vougas, 2007, Lead-Lad Relationship between Futures and Spot Markets in Greece: 1999-2001, *International Research Journal of Finance and Economics* 7, 168-174

Floros, Christos, and Dimitrios Vougas, 2006, Hedging Effectiveness in Greek Stock Index Futures Market, 1999-2001, *International Research Journal of Finance and Economics* 5, 7-18

Geweke, John, and Susan Porter-Hudak, 1983, The Estimation and Application of Long Memory Time Series Models, *Journal of Time Series Analysis* 4, 221-238

Goyal, Amit, and Alessio Saretto, 2009, Cross-Section of Option Returns and Volatility, forthcoming in the *Journal of Financial Economics*

Jackwerth, Jens, 2000, Recovering Risk Aversion from Option Prices and Realized Returns, *Review of Financial Studies* 13, 433-451

Jones, Christopher, 2006, A Nonlinear Factor Analysis of S&P 500 Index Option Returns, *Journal of Finance* 61, 2325-2363

Kenourgios, Dimitris, 2004, Price Discovery in the Athens Derivatives Exchange: Evidence for the FTSE/ASE-20 Futures Market, *Economic and Business Review* 6, 229-243

Kraus, Alan, and Robert Litzenberger, 1976, Skewness Preferences and the Valuation of Risk Assets, *Journal of Finance* 31, 1085-1100

Lin, Bing-Huei, and Jerry Wang, 2003, Systematic Skewness in Asset Pricing: An Empirical Examination of the Taiwan Stock Market, *Applied Economics* 35, 1877-1887

Lintner, John, 1965, The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets, *Review of Economics and Statistics* 47, 13-37

Liu, Xiaoquan, 2007, Returns to Trading Portfolios of FTSE 100 Index Options, *Applied Financial Economics* 17, 1211-1225

Ni, Sophie, 2006, Stock Option Returns: A Puzzle, *University of Illinois working paper*

O'Brien, Fergal, and Mark Shackleton, 2005, A Empirical Investigation of Option Returns: Overpricing and the Role of Higher Systematic Moments, *Derivatives Use, Trading and Regulation* 15, 300-330

Pettengill, Glenn, Sanjiv Sundaram, and Ike Mathur, 1995, The Conditional Relation Between Beta and Returns, *Journal of Financial and Quantitative Analysis* 30, 101-116

Reisman, Haim, and Gady Zohar, 2006, Mean Variance Analysis of Instantaneous Option Returns, *Israel Institute of Technology working paper*

Rubinstein, Mark, 1976, The Valuation of Uncertain Income Streams and the Pricing of Options, *Bell Journal of Economics* 7, 407-425

Santa-Clara, Pedro, and Alessio Saretto, 2009, Option Strategies: Good Deals and Margin Calls, forthcoming in the *Journal of Financial Markets*

Sharpe, William, 1964, Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk, *Journal of Finance* 19, 425-442

Skiadopoulos, George, 2004, The Greek Implied Volatility Index: Construction and Properties, *Applied Financial Economics* 14, 1187-1196

Vanden, Joel, 2006, Option Coskewness and Capital Asset Pricing, *Review of Financial Studies* 19, 1279-1320

Wilkens, Sascha, 2006, Option Returns Versus Asset-Pricing Theory: Evidence from the European Option Market, *Journal of Derivatives and Hedge Funds* 13, 170-176

Chapter 6
Summary and Conclusions

This thesis is motivated by the significant and increasing interest of the finance literature in the concept of volatility. The risk-return relationship constitutes one of relatively few 'axioms' in finance and states that a rational investor allocates resources into alterative investments in an attempt to maximize the overall position's expected return for a given level of risk or, alternatively, to minimize the position's risk exposure for a given level of expected return. This proposition was illustrated by Markowitz's efficient frontier which consists of all optimal portfolios that satisfy the risk-return relationship, and it has led to the development of the *Capital Asset Pricing Model*, one of the most popular and heavily discussed pricing models among academic researchers and finance practitioners alike.

The attention that volatility has received stems from it being directly related to the concept of risk, since the riskiness of an investment is commonly measured by the volatility of its returns. The accurate measurement and forecasting of returns volatility is, thus, an issue with significant implications for a variety of subjects within finance, such as portfolio management, derivatives trading, and risk management. The majority of studies that refer to volatility measurement and prediction generally fall under two main frameworks which are based on significantly different lines of thought. More specifically, the first framework assumes that the conditional variance of returns does not depend on exogenous variables and that it can be modelled by fitting *Autoregressive Conditional Heteroscedasticity* models on historical returns, while the second one relies on the prices of options to infer the market's expectation about the future volatility of the underlying asset's returns.

The present dissertation consists of four empirical essays on the two methodologies for modelling volatility that were mentioned above. These essays were presented in Chapters 2 to 5 and they examined specific research questions that relate to historical volatility models and to option-implied volatility estimates without, however, attempting to compare the respective performance of the two frameworks in measuring and forecasting volatility. Given that each empirical essay addresses a different issue and that it can be read separately rather than sequentially in the overall thesis, a set of independent conclusions was reached with respect to asymmetric historical volatility models, the informational content of implied volatility and the efficiency of options markets.

In summary, asymmetric GARCH models are examined in the case of the index as well as for individual stocks, focusing on the commonly observed higher degree of asymmetry in index volatility compared to that of stock volatilities. The dynamics of the average realized correlation among the index's constituents are found to explain part of this discrepancy, suggesting that changes in diversification opportunities have a significant impact on the asymmetry in the conditional variance of a value-weighted combination of stocks. Furthermore, the stylized fact of individual stock volatilities being less asymmetric than index volatility is shown to be easily accommodated by the 'down-market' hypothesis where the conditional variances of both asset classes respond asymmetrically to market-level innovations rather than to idiosyncratic ones. Regarding option-implied information, it is found that the implied variance of the spot exchange rate goes some way into explaining the time-varying risk-premium in currency markets and that incorporating this term in forward unbiasedness regressions results in a significant improvement of the validity of the *Uncovered Interest Parity*. Finally, it appears that the efficiency of the emerging options market in Greece is comparable to that of the developed US and UK markets in the sense that, despite higher transaction costs and thinner trading, the extent of mispricing and the opportunity for abnormal returns is not more pronounced compared to its developed counterparts.

Chapter 2, in particular, examines whether the 'diversification' hypothesis can account for the widely reported empirical finding of index volatility being more asymmetric than the volatilities of its individual components. The analysis is based on the notion that, since index variance can be decomposed to the weighted sum of individual stock variances and to the sum of cross-correlations, if asymmetric index volatility does not stem from asymmetric individual volatilities it could be the case that changes in correlations drive this conditional negative relationship. Preliminary evidence is presented of an asymmetric negative co-movement between the index's average realized correlation and past signed index returns, similarly to the asymmetric co-movement between the conditional variance and past returns. Changes in the average realized correlation are also shown to cause changes (in the Granger sense) in the realized volatility of the index, supporting the 'diversification' hypothesis. More importantly, though, estimating an extended GJR specification that includes conditional changes in the index's average correlation as an additional regressor results in a less pronounced 'asymmetry effect' in index returns. Overall, the dynamics of the average realized correlation among the index's constituents appear to be significantly correlated with the index's conditional variance and to absorb some of the explanatory power of lagged returns' magnitude

The third Chapter attempts to explain the above discrepancy in the degrees of volatility asymmetry exhibited by indices and by individual stocks from the perspective of the 'down-market' effect which suggests that returns' volatility responds asymmetrically to systematic news rather than to idiosyncratic ones. It is found that the conditional variances of individual stocks are in fact asymmetric with respect to lagged signed market returns and that the extent of asymmetry is significantly higher than the one documented with respect to firm-specific returns. Moreover, the asymmetric response of individual stock volatilities to market-level innovations is not necessarily less pronounced compared to volatility asymmetry in index returns, indicating that the 'volatility asymmetry phenomenon' is a common characteristic of both asset classes within the 'down-market' explanation.

The next two essays, presented in Chapters 4 and 5, shift the focus from historical volatility models to option-implied information. More specifically, Chapter 4 uses option-implied variance in order to re-examine the forward premium puzzle, i.e. the commonly reported finding that regressions of spot exchange rates on forward rates systematically fail to produce the unity slopes that *Uncovered Interest Parity* predicts. When the *Jensen Inequality Term* (JIT) of the spot rate's future variance, proxied by the option-implied variance, is included into the forward unbiasedness specification, the resulting proportion of slope coefficients that do not reject the UIP experiences a threefold increase compared to the standard specification. Moreover, this significant improvement in the number of rolling regressions that do not reject theoretical predictions is found to be robust to different estimation techniques for extracting the implied variance from prices of options written on foreign exchange. It is concluded, therefore, that previous studies' findings of the JIT variance having an insignificant impact on explaining deviations from the UIP could potentially be attributed to a problematic proxy for the future variance of the exchange rate rather than constituting a fundamental characteristic of currency markets.

Finally, using option prices to extract the market's expectation of the underlying's future distribution assumes the informational efficiency of the options market. Given that the efficiency of traditional options markets, such as those of the US and the UK, has already been examined in previous research, the fifth Chapter focuses on the comparable efficiency of emerging options markets. The data used refer to the Athens Derivatives Exchange in Greece and, similarly to previous studies, the efficiency of the market is evaluated from the perspective of option returns in excess of their underlying risk. Based on two commonly used criteria to measure abnormal returns, namely CAPM alphas and returns of delta and/or vega neutral straddles, it is found that the emerging options market in Athens is not characterized by a higher degree of mispricing compared to its developed counterparts. These results are consistent with the fact that most of the trading volume in the Athens exchange is attributed to large international investors with experience in more developed markets so that deviations of quoted option prices from their 'true' values are minimized. This finding is, however, in contrast to the hypothesis that higher transaction costs and thinner trading in Athens are likely to widen the no-arbitrage band and, therefore, be associated with more pronounced mispricing.

References

Andersen, Torben, and Tim Bollerslev, 1997a, Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns, *Journal of Finance* 52, 975-1005

Andersen, Torben, and Tim Bollerslev, 1997b, Intraday Periodicity and Volatility Persistence in Financial Markets, Journal *of Empirical Finance* 4, 115-158

Andersen, Torben, Tim Bollerslev, Francis Diebold, and Heiko Ebens, 2001a, The distribution of Realized Stock Return Volatility, *Journal of Financial Economics* 61, 43-76

Andersen, Torben, Tim Bollerslev, Francis Diebold, and Paul Labys, 2001b, Exchange Rate Returns Standardized by Realized Volatility are (nearly) Gaussian, *Journal of the American Statistical Association* 96, 42-55

Bacchetta, Philippe, and Eric van Wincoop, 2007, Random Walk Expectations and the Forward Discount Puzzle, *American Economic Review, Papers and Proceedings* 97, 346-350

Bacchetta, Philippe, and Eric van Wincoop, 2008, Incomplete Information Processing: A Solution to the Forward Discount Puzzle, *Study Center Gerzensee working paper*

Backus, David, Allan Gregory, and Chris Telmer, 1993, Accounting for Forward Rates in Markets for Foreign Currency, *Journal of Finance* 48, 1887-1908

Bahmani-Oskooee, Mohsen, and Satya Das, 1985, Transaction Costs and the Interest Parity Theorem, *Journal of Political Economy* 93, 793-799

Baillie, Richard, and Tim Bollerslev, 1990, A Multivariate Generalized GARCH Approach to Modeling Risk Premia in Forward Foreign Exchange Markets, *Journal of International Money and Finance* 9, 309-324

Baillie, Richard, and Tim Bollerslev, 1994, The Long Memory of the Forward Premium, *Journal of International Money and Finance* 13, 565-571

Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta-Hedged Gains and the Negative Market Volatility Risk Premium, *Review of Financial Studies* 16, 527-566

Bekaert, Geert, and Robert Hodrick, 1993, On Biases in the Measurement of Foreign Exchange Risk Premiums, *Journal of International Money and Finance* 12, 115-138

Bekaert, Geert, and Guojun Wu, 2000, Asymmetric Volatility and Risk in Equity Markets, *Review of Financial Studies* 13, 1-42

Bernoth, Kerstin, Juergen von Hagen, and Casper de Vries, 2007, The Forward Premium Puzzle: New Evidence from Futures Contracts, *DNB working paper*

Black, Fisher, 1976, Studies of Stock Price Volatility Changes, *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section*, 177-181

Black, Fisher, and Myron Scholes, 1973, The pricing of Options and Corporate Liabilities, *Journal of Political Economy* 81, 637-654

Blair, Bevan, Ser-Huang Poon, and Stephen Taylor, 2002, Asymmetric and Crash Effects in Stock Volatility for the S&P 100 Index and its Constituents, *Applied Financial Economics* 12, 319-329

Bliss, Robert, and Nikolaos Panigirtzoglou, 2002, Testing the Stability of Implied Probability Density Functions, *Journal of Banking and Finance* 26, 381-422

Bollerslev, Tim, 1986, Generalized Autoregressive Conditional Heteroscedasticity, *Journal of Econometrics* 31, 307-328

Bollerslev, Tim, Julia Litvinova, and George Tauchen, 2006, Leverage and Volatility Feedback Effects in High-Frequency Data, *Journal of Financial Econometrics* 4, 353-384

Boyer, Russell, and Charles Adams, 1988, Forward Premia and Risk Premia in a Simple Model of Exchange Rate Determination, *Journal of Money, Credit and Banking* 20, 633-644

Boyrie, Maria, Yong Kim, and Simon Pak, 2006, Price Risk and Bid-Ask Spreads of Currency Options, *Derivatives Use, Trading and Regulation* 12, 115-125

Britten-Jones, Mark, and Anthony Neuberger, 2000, Option Prices, Implied Price Processes, and Stochastic Volatility, *Journal of Finance* 55, 839-866

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2009, Understanding Index Option Returns, forthcoming in *Review of Financial Studies*

Buraschi, Andrea, and Jens Jackwerth, 2001, The Price of a Smile: Hedging and Spanning in Option Markets, *Review of Financial Studies* 14, 495-527

Byers, David, and David Peel, 1991, Some Evidence on the Efficiency of the Sterling-Dollar and Sterling-Franc Forward Exchange Rates in the Interwar Period, *Economics Letters* 35, 317-322

Campa, Jose, P.H. Chang, and Robert Reider, 1998, Implied Exchange Rate Distributions: Evidence from OTC Option Markets, *Journal of International Money and Finance* 17, 117-160

Campbell, John, and Ludger Hentschel, 1992, No News is Good News, *Journal of Financial Economics* 31, 281-318

Cappiello, Lorenzo, Robert Engle, and Kevin Sheppard, 2006, Asymmetric Dynamics in the Correlations of Global Equity and Bond Returns, *Journal of Financial Econometrics* 4, 537-572

Carr, Peter, and Liuren Wu, 2004, Time-Changes Levy Processes and Option Pricing, *Journal of Financial Economics* 71, 113-141

Carr, Peter, and Liuren Wu, 2009, Variance Risk Premiums, *Review of Financial Studies* 22, 1311-1341

Carverhill, Andrew, Terry Cheuk, and Sigurd Dyrting, 2006, The Smirk in the S&500 Futures Options Prices: A Linearized Factor Analysis, *working paper*

Chaboud, Alain, and Jonathan Wright, 2005, Uncovered Interest Parity, It Works, but not for Long, *Journal of International Economics* 66, 349-362

Cheng, Ai-Ru, Ronald Gallant, Chuanshu Ji, and Beom Lee, 2009, A Gaussian Approximation Scheme for Computation of Option Prices in Stochastic Volatility Models, forthcoming in the *Journal of Econometrics*

Chesney, Marc, and M. Jeanblanc, 2004, Pricing American Currency Options in an Exponential Levy Model, *Applied Mathematical Finance* 11, 207-225

Cheung, Yin-Wong, and Lilian Ng, 1992, Stock Price Dynamics and Firm Size: An Empirical Investigation, *Journal of Finance* 48, 1985-1997

Chinn, Menzie, 2006, The (Partial) Rehabilitation of the Interest Parity in the Floating Rate Era: Longer Horizons, Alternative Expectations, and Emerging Markets, *Journal of International Money and Finance* 25, 7-21

Chinn, Menzie, and Guy Meredith, 2004, Monetary Policy and Long Horizon Uncovered Interest Parity, *IMF Staff Papers* 51, 409-430

Christie, Andrew, 1982, The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects, *Journal of Financial Economics* 3, 407-432

Christie-David, Rohan, and Mukesh Chaudhry, 2001, Coskewness and Cokurtosis in Futures Markets, *Journal of Empirical Finance* 8, 55-81

Clinton, Kevin, 1988, Transactions Costs and Covered Interest Arbitrage: Theory and Evidence, *Journal of Political Economy* 96, 358-370

Corrado, Charles, and Tie Su, 1996, Skewness and Kurtosis in S&P500 Index Returns Implied by Option Prices, *Journal of Financial Research* 19, 175-192

Coval, Joshua, and Tyler Shumway, 2001, Expected Option Returns, *Journal of Finance*, 56, 983-1009

Crowder, William, 1994, Foreign Exchange Market Efficiency and Common Stochastic Trends, *Journal of International Money and Finance* 13, 551-564

Cumby, Robert, 1988, Is it Risk? Explaining Deviations from Uncovered Interest Parity, *Journal of Monetary Economics* 22, 279-299

De Jong, Frank, Joost Driessen, and Antoon Pelsser, 2004, On the Information in the Interest Rate Term Structure and Option Prices, *Review of Derivatives Research* 7, 99-127

Dittmar, Robert, 2002, Nonlinear Pricing Kernels, Kurtosis Preference, and Evidence from the Cross Section of Equity Returns, *Journal of Finance* 57, 368-403

Doran, James, and Andy Fodor, 2008, Is There Money to be Made Investing in Options? A Historical Perspective, *working paper*

Driessen, Joost, Pascal Maenhout, and Grigory Vilkov, 2008, The Price of Correlation Risk: Evidence from Equity Options, forthcoming in the *Journal of Finance*

Duan, Jin-Chuan, 1997, Augmented GARCH(p-q) Process and its Diffusion Limit, *Journal of Econometrics* 79, 97-217

Dumas, Bernard, Jeff Fleming, and Robert Whaley, 1998, Implied Volatility Functions: Empirical Tests, *Journal of Finance* 53, 2053-2106

Engel, Charles, 1996, The Forward Discount Anomaly and the Risk Premium: A Survey of Recent Evidence, *Journal of Empirical Finance* 3, 123-192

Engle, Robert, 1982, Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation, *Econometrica* 50, 987-1007

Engle, Robert, and Victor Ng, 1993, Measuring and Testing the Impact of News on Volatility, *Journal of Finance* 48, 1749-1778

Evans, Martin, 1995, Peso Problems: Their Theoretical and Empirical Implications, *Handbook of Statistics: Statistical Methods in Finance* (North-Holland, Amsterdam)

Fama, Eugene, and Kenneth French, 1992, The Cross-Section of Expected Returns, *Journal of Finance* 59, 427-465

Fama, Eugene, and James MacBeth, 1973, Risk, Return and Equilibrium: Empirical Tests, *Journal of Political Economy* 81, 607-636

Fang, Hsing, and Tsong-Yue Lai, 1997, Cokurtosis and Capital Asset Pricing, *Financial Review* 32, 293-307

Felmingham, Bruce, and Susan Leong, 2005, Parity Conditions and the Efficiency of the Australian 90- and 180-day Forward Markets, *Review of Financial Economics* 14, 127-145

Figlewski, Stephen, and Xiaozu Wang, 2000, Is the 'Leverage Effect' a Leverage Effect?, *NYU working paper*

Floros, Christos, and Dimitrios Vougas, 2007, Lead-Lad Relationship between Futures and Spot Markets in Greece: 1999-2001, *International Research Journal of Finance and Economics* 7, 168-174

Floros, Christos, and Dimitrios Vougas, 2006, Hedging Effectiveness in Greek Stock Index Futures Market, 1999-2001, *International Research Journal of Finance and Economics* 5, 7-18

French, Kenneth, G. William Schwert, and Robert Stambaugh, 1987, Expected Stock Returns and Volatility, *Journal of Financial Economics* 19, 3-30

Froot, Kenneth, and Jeffrey Frankel, 1989, Forward Discount Bias: Is it an Exchange Risk Premium, *Quarterly Journal of Economics* 104, 139-161

Garman, Mark, and Steven Kohlhagen, 1983, Foreign Currency Option Values, *Journal of International Money and Finance* 2, 231-237

Geweke, John, and Susan Porter-Hudak, 1983, The Estimation and Application of Long Memory Time Series Models, *Journal of Time Series Analysis* 4, 221-238

Glosten, Lawrence, Ravi Jagannathan, and David 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

Gourinchas, Pierre-Olivier, and Aaron Tornell, 2004, Exchange Rate Puzzles and Distorted Beliefs, *Journal of International Economics* 63, 303-333

Goyal, Amit, and Alessio Saretto, 2009, Cross-Section of Option Returns and Volatility, forthcoming in the *Journal of Financial Economics*

Han, Bing, 2007, Stochastic Volatilities and Correlations of Bond Yields, *Journal of Finance* 62, 1491-1524

Hodrick, Robert, 1987, The Empirical Evidence on the Efficiency of Forward and Futures Foreign Exchange Markets, (Harwood)

Hodrick, Robert, 1989b, U.S. International Capital Flows: Perspectives from Rational Maximizing Models, *Carnegie-Rochester Conference Series on Public Policy* 30, 231-288

Jackwerth, Jens, 2000, Recovering Risk Aversion from Option Prices and Realized Returns, *Review of Financial Studies* 13, 433-451

Jiang, George, and Yisong Tian, 2005, Model-Free Implied Volatility and its Information Content, *The Review of Financial Studies* 18, 1305-1342

Jones, Christopher, 2006, A Nonlinear Factor Analysis of S&P 500 Index Option Returns, *Journal of Finance* 61, 2325-2363

Kenourgios, Dimitris, 2004, Price Discovery in the Athens Derivatives Exchange: Evidence for the FTSE/ASE-20 Futures Market, *Economic and Business Review* 6, 229-243

Kim, Dogncheol, and Stanley Kon, 1994, Alternative Models for the Conditional Heteroscedasticity of Stock Returns, *Journal of Business* 67, 563-598

Kraus, Alan, and Robert Litzenberger, 1976, Skewness Preferences and the Valuation of Risk Assets, *Journal of Finance* 31, 1085-1100

Kroner, Kenneth, and Victor Ng, 1998, Modeling Asymmetric Comovements of Asset Returns, *Review of Financial Studies* 11, 817-844

Lamoureux, Christopher, and William Lastrapes, 1990, Persistence in Variance, Structural Change, and the GARCH Model, *Journal of Business and Economic Statistics* 8, 225-234

Lewis, Karen, 1994, Puzzles in International Financial Markets, *The Handbook of International Economics*, (North-Holland, Amsterdam)

Lin, Bing-Huei, and Jerry Wang, 2003, Systematic Skewness in Asset Pricing: An Empirical Examination of the Taiwan Stock Market, *Applied Economics* 35, 1877-1887

Lintner, John, 1965, The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets, *Review of Economics and Statistics* 47, 13-37

Liu, Xiaoquan, 2007, Returns to Trading Portfolios of FTSE 100 Index Options, *Applied Financial Economics* 17, 1211-1225

Longstaff, Francis, Pedro Santa-Clara, and Eduardo Schwartz, 2001, The Relative Valuation of Caps and Swaptions: Theory and Empirical Evidence, *Journal of Finance* 56, 2067-2109

Lopez, Jose, and Christian Walter, 2000, Is Implied Correlation Worth Calculating: Evidence from Foreign Exchange Options and Historical Data, *Journal of Derivatives* 7, 65-82

Malz, Allan, 1997, Option-Implied Probability Distributions and Currency Excess Returns, *Federal Reserve Bank of New York Staff Report* 32, November 1997

Mark, Nelson, Yangru Wu, and Weike Hai, 1997, Understanding Spot and Forward Exchange Rate Regressions, *Journal of Applied Econometrics* 12, 715-734

Maynard, Alex, and Peter Phillips, 2001, Rethinking an Old empirical Puzzle: Econometric Evidence on the Forward Discount Anomaly, *Journal of Applied Econometrics* 16, 671-708

McCallum, Bennett, 1994, A Reconsideration of the Uncovered Interest Parity Relationship, *Journal of Monetary Economics* 33, 105-132

Nelson, Daniel, 1991, Conditional Heteroscedasticity in Asset Returns: A New Approach, *Econometrica* 59, 347-370

Newey, Whitney, and Kenneth West, 1987, A Simple, Positive Semi-Definite Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703-708

Ni, Sophie, 2006, Stock Option Returns: A Puzzle,*University of Illinois working paper*

O'Brien, Fergal, and Mark Shackleton, 2005, A Empirical Investigation of Option Returns: Overpricing and the Role of Higher Systematic Moments, *Derivatives Use, Trading and Regulation* 15, 300-330

Pettengill, Glenn, Sanjiv Sundaram, and Ike Mathur, 1995, The Conditional Relation Between Beta and Returns, *Journal of Financial and Quantitative Analysis* 30, 101-116

Phillips, Peter, 1995, Robust Nonstationary Regression, *Econometric Theory* 11, 912-951

Phillips, Peter, James McFarland, and Patrick McMahon, 1996, Robust Tests of Forward Exchange Market Efficiency with Empirical Evidence from the 1920s, *Journal of Applied Econometrics* 11, 1-22

Pollet, Joshua, and Mungo Wilson, 2008, Average Correlation and Stock Market Returns, forthcoming in the *Journal of Financial Economics*

Pong, Shiuyan, Mark Shackleton, Stephen Taylor, and Xinzhong Xu, 2004, Forecasting Currency Volatility: A Comparison of Implied Volatilities and AR(FI)MA Models, *Journal of Banking and Finance* 28, 2541-2563

Reisman, Haim, and Gady Zohar, 2006, Mean Variance Analysis of Instantaneous Option Returns, *Israel Institute of Technology working paper*

Rubinstein, Mark, 1985, Nonparametric Test of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Options Classes from August 23, 1976 Throughout August 31, 1978, *Journal of Finance* 40, 455-480

Rubinstein, Mark, 2000, Comments on the 1987 Stock Market Crash: 11 Years Later, *Risks in Investment Accumulation Products of Financial Institutions*, published by the Society of Actuaries

Santa-Clara, Pedro, and Alessio Saretto, 2009, Option Strategies: Good Deals and Margin Calls, forthcoming in the *Journal of Financial Markets*

Sharpe, William, 1964, Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk, *Journal of Finance* 19, 425-442

Skiadopoulos, George, 2004, The Greek Implied Volatility Index: Construction and Properties, *Applied Financial Economics* 14, 1187-1196

Skintzi, Vasiliki, and Apostolos Refenes, 2005, Implied Correlation Index: A New Measure of Diversification, *Journal of Futures Markets* 25, 171-197

Smith, Daniel, 2007, Risk and Return in Stochastic Volatility Models: Volatility Feedback Matters!, *Simon Fraser University working paper*

Stivers, Christopher, Patrick Dennis, and Stewart Mayhew, 2006, Stock Returns, Implied Volatility Innovations and the Implied Volatility Phenomenon, *Journal of Financial and Quantitative Analysis* 41, 381-406

Tauchen, George, Harold Zhang, and Ming Liu, 1996, Volume, Volatility and Leverage: A Dynamic Analysis, *Journal of Econometrics* 74, 177-208

Taylor, Stephen, 1986, Moddeling Financial Time Series, *John Wiley & Sons, Chichester*

Taylor, Stephen, Pradeep Yadav, and Yuanyuan Zhang, 2006, The Information Content of Implied Volatilities and Model-Free Volatility Expectations: Evidence from Options Written on Individual Stocks, *Lancaster University working paper*

Taylor, Stephen, and Xinzhong Xu, 1994, The Magnitude of Implied Volatility Smiles: Theory and Empirical Evidence for Exchange Rates, *Review of Futures Markets* 13, 355-380

Vanden, Joel, 2006, Option Coskewness and Capital Asset Pricing, *Review of Financial Studies* 19, 1279-1320

Wilkens, Sascha, 2006, Option Returns Versus Asset-Pricing Theory: Evidence from the European Option Market, *Journal of Derivatives and Hedge Funds* 13, 170-176

Wu, Guojun, 2001, The Determinants of Asymmetric Volatility, *Review of Financial Studies* 14, 837-859