

University of Warwick institutional repository: http://go.warwick.ac.uk/wrap

A Thesis Submitted for the Degree of PhD at the University of Warwick

http://go.warwick.ac.uk/wrap/66718

This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it. Our policy information is available from the repository home page.



Foundations of Equity Market Leverage Effects

Marcus Alexander Ong

Submitted for the Qualification of

Doctor of Philosophy in Complexity Science with Finance

Centre for Complexity Science, University of Warwick

December 2014

THE UNIVERSITY OF WARWICK

Contents

Li	st of I	igures	14
Li	st of]	ables	16
Ac	cknow	edgments	17
De	eclara	ion	17
Su	ımma	y	18
Al	brevi	ations	19
Ι	Th	Leverage Effect in Stocks	20
1	Lite	ature Review: The Stock Level Leverage Effect	21
	1.1	Introduction	 21
	1.2	Leverage Hypothesis	 22
	1.3	Volatility Feedback Hypothesis	 28
	1.4	Retarded Volatility Hypothesis	 28
	1.5	Volume Hypothesis	 30
		1.5.1 Returns and Volume	 33
		1.5.2 Volume and Volatility	 34
		1.5.3 Behavioural Biases and Heuristics	 35
	1.6	Summary	 37

2	An]	nformation Theoretic Analysis of Stock Returns, Volatility & Trading	
	Volu	nes 3	8
	2.1	Introduction	8
	2.2	Literature Review	9
		2.2.1 Granger Causality	9
		2.2.2 Information Theory	1
	2.3	Methodology	5
		2.3.1 The Data	5
		2.3.2 Calibrating the Model	5
	2.4	Results	8
		2.4.1 Persistence of Returns, Volatility and Volumes	8
		2.4.2 Returns and Volumes	60
		2.4.3 Volume and Volatility	;3
		2.4.4 Returns and Volatility	6
	2.5	Conclusions	64
3	Retu	rn-Volume Correlation and its Behavioural Origins 6	5
	3.1	Introduction	55
	3.2	Literature Review	57
		3.2.1 Prospect Theory	57
	3.3	Model (Barberis and Xiong, 2009)	0'
		3.3.1 Simulations	'2
		3.3.2 The Effect of Expected Return and Standard Deviation 7	'3
		3.3.3 The Effect of the Utility Function	'7
	3.4	Methodology	0
		3.4.1 The Data	60
	3.5	Results	31
		3.5.1 Return-Volume Correlation	31
		3.5.2 Model Predictions	3
	3.6	Conclusions	17

II	T	ne Leverage Effect in Stock Indices	99
4	Lite	rature Review: The Index Level Leverage Effect	100
	4.1	Introduction	100
	4.2	Leverage Hypothesis	100
	4.3	Volatility Feedback Hypothesis	102
	4.4	Retarded Volatility Hypothesis	103
	4.5	Volume Hypothesis	103
	4.6	Models of Collective Behaviour	104
	4.7	Stock Correlation Hypothesis	105
		4.7.1 Empirical Evidence	105
		4.7.2 Relationship between the Index and Stock Levels	106
	4.8	Summary	107
5	An	Information Theoretic Analysis of Index Returns, Volatility and Trading	T
-	Volu	imes	, 108
	5.1	Introduction	108
	5.2	Methodology	109
		5.2.1 The Data	109
		5.2.2 Calibrating the Model	109
	5.3	Results	113
		5.3.1 Persistence of Returns, Volatility and Volumes	113
		5.3.2 Returns and Volumes	115
		5.3.3 Volume and Volatility	118
		5.3.4 Returns and Volatility	121
	5.4	Conclusions	126
Π	II	ime Variation of the Leverage Effect	127
6	Tim	e-Series Analysis of the Leverage Effect	128
	6.1	Introduction	128

	6.2	Methodology	29
		6.2.1 The Data	29
		6.2.2 What are Wavelets?	29
	6.3	Results	31
		6.3.1 Stock Level Effects	31
		6.3.2 Index Level Effects	35
		6.3.3 Index and Stock Level Effects	36
	6.4	Conclusions	40
I	/ Τ	he Leverage Effect in Stock Options 14	41
7	Lite	rature Review: The Implied Leverage Effect 1	.42
	7.1	Properties of Implied Volatility	.43
	7.2	Models	.44
	7.3	What drives the Implied Leverage Effect?	46
	7.4	Summary	47
8	Mod	lelling Implied Volatilities using Multivariate q-Gaussian Distributions 1	48
	8.1	Introduction	48
	8.2	q-Gaussian Distributions	49
	8.3	Methodology	53
		8.3.1 The Data	53
		8.3.2 Calibrating the Model	54
	8.4	Results	55
		8.4.1 Fitting the q-Gaussian Distribution	55
		8.4.2 The Implied Leverage Effect	60
		8.4.3 Stressing the Implied Volatility Surface	61
		8.4.4 Virtual Data Generation	62
	8.5	Conclusions	64

V	Co	onclusi	ions	165
Bi	bliogi	raphy		169
Aj	ppen	dices		188
A	An]	Informa	ntion Theoretic Analysis of Returns, Volatility and Trading V	′ol-
	ume	s at the	Weekly Frequency	188
	A.1	Stock	Results	. 189
		A.1.1	Calibrating the Model	. 189
		A.1.2	Persistence of Returns, Volatility and Volumes	. 190
		A.1.3	Returns and Volume	. 191
		A.1.4	Volume and Volatility	. 193
	A.2	Index	Results	. 195
		A.2.1	Calibrating the Model	. 195
		A.2.2	Persistence of Returns, Volatility and Volumes	. 196
		A.2.3	Returns and Volume	. 197
		A.2.4	Volume and Volatility	. 199
		A.2.5	Returns and Volatility	. 201
B	Info	rmatior	n Theoretic Analysis of Returns, Volatility and Trading Volum	nes
	by I	ndex		203
	B .1	The Da	ata	. 203
	B.2	Result	s by Index	. 205
		B.2.1	Returns and Volume	. 205
		B.2.2	Volume and Volatility	. 207
		B.2.3	Returns and Volatility	. 209
C	An]	Informa	ation Theoretic Analysis of Returns, Volatility and Trading V	′ol-
	ume	s for Er	nerging Market Indices	213
		C.0.4	The Data	. 213

	C.1.4	Returns and Volatility	219
	C.1.3	Volume and Volatility	217
	C.1.2	Returns and Volume	215
	C.1.1	Persistence of Returns, Volatility and Volumes	214
C.1	Emerg	ing Markets Results	214

List of Figures

2.1	Partial Mutual Information (<i>PMI</i>) Diagram	42
2.2	Mutual Information Calibration for Stocks at the Daily Frequency	47
2.3	Auto-Mutual Information and Auto-Covariance for Stocks at the Daily	
	Frequency	49
2.4	Cross-Mutual Information and Cross-Covariance Functions for Stock Re-	
	turns and Volume at the Daily Frequency	51
2.5	Partial Cross-Mutual Information and Cross-Transfer Entropy Functions	
	for Stock Returns and Volumes at the Daily Frequency	52
2.6	Cross-Mutual Information and Cross-Covariance Functions for Stock Volatil-	
	ity and Volume at the Daily Frequency	54
2.7	Partial Cross-Mutual Information and Cross-Transfer Entropy Functions	
	for Stock Volatility and Volumes at the Daily Frequency	55
2.8	Cross-Mutual Information and Cross-Covariance Functions for Stock Re-	
	turns and Volatility at the Daily Frequency	57
2.9	Partial Cross-Mutual Information and Cross-Transfer Entropy Functions	
	for Stock Returns and Volatility at the Daily Frequency	58
2.10	Cross-Mutual Information and Cross-Covariance Functions for Stock Re-	
	turns and Volatility at the Weekly Frequency	60
2.11	Partial Cross-Mutual Information and Cross-Transfer Entropy Functions	
	for Stock Returns and Volatility at the Weekly Frequency	61
3.1	Prospect Theory Utility Curve	69

3.2	Return-Volume Correlation calculated using the Barberis and Xiong (2009)
	model for Linear and Prospect Theory Preferences
3.3	Trading Behaviour calculated using the Barberis and Xiong (2009) Model
	for Linear and Prospect Theory Preferences
3.4	Prospect Theory Utility Curve for Varying Curvature
3.5	Effect of Varying Curvature and Loss Aversion on Return-Volume Corre-
	lation in the Barberis and Xiong (2009) Model
3.6	Density Plots of Return-Volume Correlation for Developed and Emerging
	Market Stocks
5.1	Mutual Information Calibration for Indices at the Daily Frequency 112
5.2	Auto-Mutual Information and Auto-Covariance Functions for Indices at
	the Daily Frequency
5.3	Cross-Mutual Information and Cross-Covariance Functions for Index Re-
	turns and Volume at the Daily Frequency
5.4	Partial Cross-Mutual Information and Cross-Transfer Entropy Functions
	for Index Returns and Volume at the Daily Frequency
5.5	Cross-Mutual Information and Cross-Covariance Functions for Index Volatil-
	ity and Volume at the Daily Frequency
5.6	Partial Cross-Mutual Information and Cross-Transfer Entropy Functions
	for Index Volatility and Volume at the Daily Frequency
5.7	Cross-Mutual Information and Cross-Covariance Functions for Index Re-
	turns and Volatility at the Daily Frequency
5.8	Partial Cross-Mutual Information and Cross-Transfer Entropy Functions
	for Index Returns and Volatility at the Daily Frequency
6.1	Time Variation of Return-Volume Correlation and the Sharpe Ratio 133
6.2	Time Variation of the Stock Level Leverage Effect and Return-Volume
	Correlation

6.3	Time Variation of the Index Level Leverage Effect and Return-Volume
	Correlation
6.4	Time Variation of the Index and Stock Level Leverage Effects 136
6.5	Time Variation of the Index Level Leverage Effect and Correlation Asym-
	metry
6.6	Time Variation of the Stock Level Leverage Effect and Correlation Asym-
	metry
6.7	Time Variation of the Transfer Entropy from the Index to Stock Level 139
8.1	Univariate q-Gaussian Distribution
8.2	Time Variation of the Goldman Sachs Stock Price and At-The-Money
	Implied Volatility
8.3	Time Variation of Goldman Sachs Stock Returns and Changes in At-The-
	Money Implied Volatility
8.4	MCMC Calibration for Multivariate q-Gaussian
8.5	Univariate q-Gaussian Fits of Returns and Implied Volatility Changes for
	Goldman Sachs
8.6	Bivariate q-Gaussian Fits of Returns and Implied Volatility Changes for
	Goldman Sachs
8.7	Fitted q-Gaussian Parameters for S&P500 Stocks
8.8	Histogram of the Akaike Information Criterion for S&P500 stocks 158
8.9	Time variation of q and β for Goldman Sachs Stock $\ldots \ldots \ldots$
8.10	Implied Leverage Effect for Goldman Sachs Stock
8.11	Decay of the Instantaneous Skew for Goldman Sachs Stock
8.12	Goldman Sachs Autocorrelations
8.13	Virtual Data Generation
A.1	Mutual Information Calibration for Stocks at the Weekly Frequency 189
A.2	Auto-Mutual Information and Auto-Covariance Functions for Stocks at
	the Weekly Frequency

A.3	Cross-Mutual Information and Cross-Covariance Functions for Stock Re-
	turns and Volume at the Weekly Frequency
A.4	Cross-Partial Mutual Information and Cross-Transfer Entropy Functions
	for Stock Returns and Volume at the Weekly Frequency
A.5	Cross-Mutual Information and Cross-Covariance Functions for Stock Volatil-
	ity and Volume at the Weekly Frequency
A.6	Cross-Partial Mutual Information and Cross-Transfer Entropy Functions
	for Stock Volatility and Volume at the Weekly Frequency
A.7	Mutual Information Calibration for Indices at the Weekly Frequency 195
A.8	Auto-Mutual Information and Auto-Covariance Functions for Indices at
	the Weekly Frequency
A.9	Cross-Mutual Information and Cross-Covariance Functions for Index Re-
	turns and Volume at the Weekly Frequency
A.10	Cross-Partial Mutual Information and Cross-Transfer Entropy Functions
	for Index Returns and Volume at the Weekly Frequency
A.11	Cross-Mutual Information and Cross-Covariance Functions for Stock Volatil-
	ity and Volume at the Weekly Frequency
A.12	Cross-Partial Mutual Information and Cross-Transfer Entropy Functions
	for Index Volatility and Volume at the Weekly Frequency
A.13	Cross-Mutual Information and Cross-Covariance Functions for Index Re-
	turns and Volatility at the Weekly Frequency
A.14	Cross-Partial Mutual Information and Cross-Transfer Entropy Functions
	for Index Returns and Volatility at the Weekly Frequency
C.1	Auto-Mutual Information and Auto-Covariance Functions for Emerging
	Market Indices at the Daily Frequency
C.2	Cross-Mutual Information and Cross-Covariance Functions for Emerging
	Market Index Returns and Volume at the Daily Frequency
C.3	Cross-Partial Mutual Information and Cross-Transfer Entropy Functions
	for Emerging Market Index Returns and Volume at the Daily Frequency . 216

C.4	Cross-Mutual Information and Cross-Covariance Functions for Emerging	
	Market Index Volatility and Volume at the Daily Frequency	217
C.5	Cross-Partial Mutual Information and Cross-Transfer Entropy Functions	
	for Emerging Market Index Volatility and Volume at the Daily Frequency	218
C.6	Cross-Mutual Information and Cross-Covariance Functions for Emerging	
	Market Index Returns and Volatility at the Daily Frequency	219
C.7	Cross-Partial Mutual Information and Cross-Transfer Entropy Functions	
	for Emerging Market Index Returns and Volatility at the Daily Frequency	220
D.1	Univariate q-Gaussian Fits of Returns and Implied Volatility Changes for	
	US Stocks	223
D.2	Akaike Information Criterion difference between the Gaussian and q-	
	Gaussian Models	224

List of Tables

2.1	Effective Mutual Information for the Stock Level Leverage Effect Con-
	trolling for Trading Volumes and Index Returns
3.1	Regression for Return-Volume Correlation and Sharpe Ratio for Devel-
	oped Market Indices
3.2	Regression for Return-Volume Correlation and Sharpe Ratio for Emerg-
	ing Market Indices
3.3	Panel Regression for Return-Volume Correlation and Sharpe Ratio 89
3.4	Regression for Return-Volume Correlation and Trading Strategy (S) for
	S&P500 Institutional Trading
3.5	Return-Volume Correlation Grouped into Developed and Emerging Markets 93
3.6	Regression for Return-Volume Correlation and Loss-Aversion (λ) for
	S&P500 Stocks
4.1	Leverage Effect at the Stock and Index Levels from Figlewski and Wang
	(2000)
5.1	Mutual Information for the Index Level Leverage Effect Controlling for
	Trading Volumes and Correlation Asymmetry
B .1	Summary Statistics by Index
B.2	Summary of the Stock Level Return-Volume Relation by Index 205
B.3	Summary of the Stock Level Volume-Volatility Relation by Index 207
B.4	Summary of the Stock Level Return-Volatility Relation by Index 209

B.5	Effective Mutual Information for the Stock Level Leverage Effect by In-	
	dex Controlling for Trading Volumes and Index Returns	211

Acknowledgments

This thesis is dedicated to the memory of my grandfather Alan Stacey who inspired and encouraged my interest in science and mathematics. It is also dedicated to my mother Lesley Ong and my Grandmother Irene Stacey without whose support and encouragement this work would not have been possible. I would like to thank my wonderful family for whom I am eternally grateful.

I also would like to thank my supervisors Robin Ball, Roman Kohzan and Anthony Neuberger for their guidance and the EPSRC for funding this research.

Declaration

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy in Complexity Science with Finance. I composed the thesis and have not submitted any of the research in previous applications for any degree.

The research presented, including data generated and data analysis, was carried out by myself except in the cases outlined below:

1) The results in Table 3.3 were produced by my supervisor Roman Kohzan as I was unfamiliar with the statistical analysis system, STATA, which includes fixed effect controlled regressions as inbuilt functions.

2) The research in Chapter 8 and Appendix D was conducted in collaboration with my colleague Daniel Sprague. The research was instigated, designed and prototyped by myself with Daniel advising on the implementation and calibration of the Markov Chain Monte Carlo (MCMC). Daniel also implemented a faster version of the algorithm using the 'PyMC' Python library which was used to produce the figures in Chapter 8 and the Akakie Information Criterion (AIC) method for comparing unnormalised multivariate probability distributions.

The research in Chapters 2, 3 and 8 have been submitted for publication in Finance journals and I am currently awaiting feedback.

Summary

This thesis examines the Leverage Effect in stocks, stock indices and stock options. The Leverage Effect refers to the observed negative correlation between an asset's return and its volatility. Part I presents an examination of the Leverage Effect at the stock level. The research provides the first investigation of stock returns, volatility and trading volumes from an information theoretic perspective. It finds support for trading volumes as an explanation for the stock level Leverage Effect and shows that index returns are also an important factor. It also analyses how trading behaviour is influenced by an investor's risk preference and how this relates to return-volume correlation. Predictions of an analytical model of trading behaviour are verified empirically using a range of stocks and institutional trades in S&P500 stocks. Part II examines the Leverage Effect at the index level. The research supports previous findings that the Leverage Effect is far larger at the index level and decays more quickly. Again using an information theoretic analysis, it shows that it is driven by a combination of trading volumes and an asymmetric relationship between index returns and stock return correlations. Part III examines the time variation of the Leverage Effect at the stock and index levels. It shows that they are both time dependent and discusses the relationship between the stock and index levels. It also documents changes in market behaviour since the 2008 financial crisis. Part IV examines the Leverage Effect in stock options by developing a descriptive statistical model of implied volatility using multivariate q-Gaussian distributions. This is the first research to show that implied volatility can be modelled using q-Gaussian distributions and provides a tool for trading and risk management. It also shows how the multivariate q-Gaussian distribution could be used to generate virtual data for scenario testing and option pricing using a simple Markov Chain or Auto-regressive process. Finally PartV presents the conclusions of the thesis and avenues for future research.

Abbreviations

- ATMF At-The-Money-Forward
- **EMI** Effective Mutual Information
- IV Implied Volatility
- **IVS** Implied Volatility Surface
- MCMC Makov-Chain Monte Carlo
- MI Mutual Information
- MLE Maximum Likelihood Estimation
- **PMI** Partial Mutual Information
- PT Prospect Theory
- **TE Transfer Entropy**

Part I

The Leverage Effect in Stocks

Chapter 1

Literature Review: The Stock Level Leverage Effect

1.1 Introduction

The Leverage Effect refers to the negative correlation between an asset's return and its volatility. The effect was first documented by Black (1976) and it is to this work that the effect lends its name. The effect has subsequently been evidenced by many authors and is now regarded as a stylised fact (Cont, 2001). The effect may also be referred to as return-volatility asymmetry or simply volatility asymmetry. The effect may seem obvious, since volatility is defined as a relative quantity; it naturally increases when the price decreases simply because the price appears in the denominator. However, despite nearly four decades since its identification there remains significant disagreement as to both the cause and the direction of causation i.e. does a drop in an asset's price trigger greater volatility or does an increase in volatility lead to a fall in an asset's price? Many authors have also suggested that the effect is contemporaneous. The size of the Leverage Effect also appears to vary between different types of assets. For instance, it appears far stronger at the index level than at the stock level. Different models have been proposed to explain the Leverage Effect but with varying success.

In this chapter I review the proposed models for the Leverage Effect at the stock

level. The 'Leverage Hypothesis' (Section 1.2), proposed by Black (1976), speculates that the Leverage Effect is the result of changes in a firm's financial leverage which leads to changes in expectations over a firm's risk and hence its volatility. The 'Volatility Feedback Hypothesis' (Section 1.3) has been considered by numerous authors, including Campbell and Hentschel (1992), Bekaert and Wu (2000) and Aydemir et al. (2006). This conjectures that when news enters the market there is an increase in market volatility. As a result investors' expectations of future returns increase and the price consequently decreases. The 'Retarded Volatility Hypothesis' (Section 1.4) developed by Bouchaud et al. (2001). This proposes that the Leverage Effect has no economic significance but is the result of a lagged response to price changes. Finally, the 'Volume Hypothesis' which conjectures that the Leverage Effect can be explained by investors' trading behaviour (Section 1.5). I then present a summary of the findings in Section 1.6.

1.2 Leverage Hypothesis

The negative correlation between a stock's return and its volatility was first identified by Black (1976). He proposed that a drop in the value of a firm's equity will cause a negative return and thus increase the stock's leverage (i.e. its debt/equity ratio), this increase in leverage will, in turn, lead to a rise in the volatility of the stock. A similar effect may arise even if the firm has almost no debt because of 'operating leverage'. These are fixed costs that cannot be eliminated, at least in the short run, hence when expected revenues fall, profit margins decline as well. The effect therefore became known as the Leverage Effect. I will now summarise the empirical evidence.

Black (1976) - studied 30 US Stocks (1964-1975), at the monthly frequency, using the linear regression:

$$\frac{\sigma_{t+1} - \sigma_t}{\sigma_t} = a_0 + a_1 r_t + \varepsilon_{t+1} \tag{1.1}$$

where r_t is the stock return at time, t, and σ_t is the volatility at time, t (calculated as the square root of the sum of daily squared returns).

Although, Black did not report detailed results of his regressions he found that $a_1 < 0$ and usually $a_1 < -1$.

Christie (1982) - studied 379 US Stocks (1962 - 1978), at the quarterly frequency, using the linear regression:

$$\ln\left(\frac{\sigma_{t+1}}{\sigma_t}\right) = a_0 + a_1 r_t + \varepsilon_{t+1} \tag{1.2}$$

where r_t is the daily returns at time, t, and σ_t is the volatility at time, t.

Christie found that $E(a_1) = -0.23$. This was supported by Figlewski and Wang (2000) at the monthly and quarterly frequency using 100 US Stocks (1977-1996). Christie (1982) also found that a_1 was strongly negatively correlated with financial leverage (Debt/Equity Ratio). Thus he concluded that leverage was the dominant factor but probably not the sole factor in determining a_1 .

Cheung and Ng (1992) - studied 251 US Stocks (1962-1989), at the daily frequency, using an exponential Generalised Auto-Regressive Conditional Heteroskedacity (eGARCH) model:

$$\ln(\sigma_t) = a_0 + a_1 r_{t-1} + a_2 |r_{t-1}| + a_3 \ln(\sigma_{t-1})$$
(1.3)

where the conditional volatility, σ_t , depends upon the lagged volatility, σ_{t-1} , the lagged absolute returns, $|r_{t-1}|$ and the lagged returns, r_{t-1} .

The shock to the stock return, ε_t , on day, t, is given by

$$\varepsilon_t = \sigma_t r_t \tag{1.4}$$

Cheung and Ng (1992) found that $a_1 < 0$ for over 95% of the firms and confirmed the negative correlation between a_1 and financial leverage. They also found a strong positive correlation between a_1 and firm size.

Duffee (1995) - studied 2,500 US Stocks (he separates continuously traded firms) (1977-1991), at the daily and monthly frequency, using a linear regression similar to Christie (1982) and the contemporaneous relationships:

$$\ln\left(\sigma_{t}\right) = b_{0} + b_{1}r_{t} + \varepsilon_{t,1} \tag{1.5}$$

$$\ln(\sigma_{t+1}) = b_0 + b_2 r_t + \varepsilon_{t+1,2} \tag{1.6}$$

the usual lagged stock return coefficient is the difference, $a_1 \equiv b_2 - b_1$.

Duffee made several findings 1) for a 'typical firm', b_1 is strongly positive whilst the sign of b_2 is dependent upon the frequency over which the relations are estimated; it is positive at the daily frequency and negative at the monthly frequency. However, regardless of the sign of b_2 , it is always the case that $b_1 > b_2$, implying that $a_1 < 0$. 2) The negative relationship between a_1 and financial leverage is confirmed only with monthly data, for the subset of continuously traded firms. For a larger sample of firms, without survivorship bias, the correlation between turns positive. 3) a_1 is positively correlated with size but b_1 and b_2 are negatively correlated with size. Therefore, the positive correlation of a_1 with size is a consequence of the fact that the size effect in b_1 is stronger than in b_2 . This conflicts with the results of Cheung and Ng (1992) but Duffee suggests that this may be because their eGARCH model produces far smoother estimates of stock volatility than absolute daily returns; more akin to the monthly data. 4) The contemporaneous relation between returns and volatility is much greater for firms that are eventually delisted. Duffee concludes that the standard interpretation for $a_1 < 0$, that an increase in r_t corresponds to a decrease in σ_{t+1} , is incorrect. Rather the effect is mainly due to the positive contemporaneous relationship between returns and volatility. He argues that since the 'Leverage Hypothesis' has no implications for the strength of the contemporaneous relationship, there must be another reason for at least part of the correlation between firm debt/equity ratios and the regression coefficient, a_1 . Using longer period lagged returns, Hasanhodzic and Lo (2011) and Daouk and Ng (2011) claim to show that this relationship is not contemporaneous. However, this simply shows that the effect is persistent as shown by Bouchaud et al. (2001).

Hasanhodzic and Lo (2011) - studied >700 US Stocks (they separated all-equity

and all-debt financed firms) (1972-2008), at the daily and monthly frequency, using the regressions of Black (1976), Christie (1982) and Duffee (1995).

They found that all-equity financed firms generally have a more negative value for a_1 than the all-debt financed firms. Hence they concluded that the Leverage Effect was not the result of leverage and instead favour the 'Volatility Feedback Hypothesis' (Section 1.3) or 'Behavioural Models' (Section 1.5).

Figlewski and Wang (2000) - studied 100 US Stocks (1977-1996), at the monthly and quarterly frequency, testing several implications of the 'Leverage Hypothesis':

1) A positive return should decrease leverage and volatility by the same amount as it increases it for a negative return of the same size. To examine this they used the following regression to separate up and down markets effects:

$$\ln\left(\frac{\sigma_{t+1}}{\sigma_t}\right) = b_0 + b_1 r_t + b_2 f_t r_t \tag{1.7}$$

where f_t is a dummy variable ($f_t = 1$ for $r_t < 0$ and $f_t = 0$ otherwise). The leverage effect is measured by b_1 in an up market and ($b_1 + b_2$) in a down market. They found that b_2 is far more negative than b_1 and is statistically significant. This implies that there is a strong impact on volatility when stock prices fall and a much weaker effect or even positive when they rise. This led Figlewski and Wang (2000) to propose a new label for the Leverage Effect,: the 'down-market effect'. Interestingly, b_1 is generally positive which implies that volatility increases when the stock price increases. However, Daouk and Ng (2011) say this is because they do not account for the Auto-Regressive Conditional Heteroskedacity (ARCH) effect. They state that "When positive returns are large, ARCH effect pushes volatility higher while the Leverage Effect pushes volatility lower. Since ARCH is the dominant effect, volatility increases. It does not mean that the Leverage Effect is not important".

2) The amount of leverage in the firm's financial structure should determine the volatility and not the change in leverage. Hence a permanent change in leverage should produce a permanent change in volatility. To examine if the change in volatility related to price changes is permanent or dies out over time they used the following regression:

$$\ln\left(\frac{\sigma_{t+1}}{\sigma_t}\right) = b_0 + b_1 r_t + b_2 r_{t-1} + b_3 r_{t-2}$$
(1.8)

If the Leverage Effect was entirely due to the actual change in firm leverage associated with a change in the stock price, the estimates for b_1 , b_2 , and b_3 should be of approximately equal size. By contrast, a spike in volatility resulting from a non-permanent factor, such as a burst of 'irrational exuberance' in the market, may be expected to die out once the market calms down. This can then be modified to allow for an asymmetry between up and down moves:

$$\ln\left(\frac{\sigma_{t+1}}{\sigma_t}\right) = b_0 + b_1 r_t + b_2 r_{t-1} + b_3 r_{t-2} + b_4 f_t r_t + b_5 f_{t-1} r_{t-1} + b_6 f_{t-2} r_{t-2}$$
(1.9)

They found that the effect tends to die out over time.

3) The mechanism through which leverage changes should be unimportant i.e. altering the level of debt or shares outstanding should have the same effect on volatility as stock price changes. In order to conduct this analysis they analysed the following regression:

$$\ln\left(\frac{\sigma_{t+1}}{\sigma_{t-3}}\right) = b_0 + b_1 k \triangle D + b_2 k \triangle N + b_3 k r_t + b_4 k f_t \triangle D + b_5 k f_t \triangle N + b_6 k f_t r_t \quad (1.10)$$

where ΔD is the quarterly change in the log of the book value of the firm debt, ΔN is the change in the log of the number of shares outstanding and k is a multiplicative constant. They found that changes in stock prices in the market seem to account for the entire Leverage Effect with combined influence $(b_3 + b_6)$ being greater than 1. In fact, if there is any impact from share or debt issuance, these results suggest that the effect is positive. This finding is in accordance with the commonly accepted behavioural interpretation. This states that firms issue equity when it is overvalued (decrease leverage) and retire equity when it is undervalued (increase leverage). The former is seen as negative and the price tends to fall whilst the latter is positive and the price tends to rally. This is in accordance with the findings of Pontiff and Woodgate (2008). They found that

during the post-1970 time period, share issuance is strongly correlated with future returns ranging from one month to three years. In the case of annual share issuance, statistical significance is greater than the previously documented predictability attributed to bookto-market, size, and momentum.

Daouk and Ng (2011) - studied >1400 US Stocks (1987-2003), at the monthly frequency, using several methods for unlevering stocks using the market price of debt. They found that the Leverage Effect was generally not statistically significant once the stocks were unlevered. They also found that effect of financial leverage is negative and significant at the 1% level but there does not seem to be any significant relationship between operating leverage and the Leverage Effect.

Kogan (2004) found that firm investment activity and firm characteristics, particularly the book-to-market ratio might lead to the Leverage Effect. He developed this hypothesis through the development of a production economy model in which real investment is irreversible and subject to convex adjustment costs. However, Daouk and Ng (2011) found no observable relation between the Leverage Effect and beta or book-to-market value at the firm-level.

Hens and Steude (2009) - studied 24 students in an experimental stock market with an artificial stock that had no leverage. The stock paid a dividend (the true stochastic nature of which was unknown to participants) at the end of each period, through a double auction mechanism. The experimental design was based upon the consumption based asset pricing model of Breeden (1979).

The Leverage Effect naturally emerges in the four experimental markets. Using $C_{rv}(\tau) = Corr(r_t, r_{t+\tau}^2)$, they discovered $C_{rv}(1) = -0.26(0.0158), -0.44(0.0000), -0.52(0.0000)$ and -0.39(0.0026); with associated p-values shown in brackets. C_{rv} is negative and statistically significant at the 5% level in all four markets. However, only two of the series exhibit leverage with two or three lags and these are not statistically significant. Their three control experiments also demonstrated leverage although not to the same degree $C_{rv}(1) = -0.03; -0.24; -0.04$ - this may be because they were run for shorter periods which shows that the exact generating process has no influence on the Leverage Effect.

1.3 Volatility Feedback Hypothesis

Black termed this the 'reverse causation effect' and it refers to the causal relationship from volatility changes to stock returns; it is also known as the time-varying risk premium hypothesis. This hypothesis speculates that an increase in volatility implies an increase in the future required expected return of the stock, hence a decline in the current stock price. In addition, due to the persistent nature of volatility (large realizations of either good or bad news increase both current and future volatility), a feedback loop is created: the increased current volatility raises expected future volatility and therefore, expected future returns, causing stock prices to fall now.

Smith and Yamagata (2011) – studied 242 US Stocks (1973-2007), at the monthly frequency using panel vector auto-regression. They found that volatility feedback effects are present at the firm level due to market and firm effects with market effects being significantly stronger. They also identify significant leverage effects which persist for three to four months. Aydemir et al. (2006) developed a general equilibrium model that generates macroeconomic conditions that led to realistic dynamics for the risk-less rate and the market price of risk. In this model the economy is driven by the counter-cyclical risk-aversion caused by external habit formation in the representative agent's preferences. They conclude that at the stock level, financial leverage is an important driver of stock volatility dynamics but in bad times they are still driven by the market conditions. Daouk and Ng (2011) claim that this is in line with their empirical investigations. Most of the other research on the volatility feedback hypothesis has focussed on stock market indices or portfolios of stocks. Bekaert and Wu (2000) aggregated individuals stocks into small portfolios but, as Daouk and Ng (2011) state, this aggregation process diminishes the importance of the Leverage Effect at the firm-level.

1.4 Retarded Volatility Hypothesis

Bouchaud et al. (2001) argue that the stock level Leverage Effect can be interpreted within a 'retarded' model where the amplitude of the price changes does not follow the instan-

taneous price level but rather absolute price changes are related to an average level of the past price. This reflects the lag with which market operators change their behaviour (order volumes, bid-ask spreads, transaction costs etc.) when the price evolves. Bouchaud et al. (2001) state that "although it is true that on the long run, price increments tend to be proportional to prices themselves, this is not reasonable at short time scales. Locally, prices evolve in discrete steps (ticks), following buy or sell orders that can only be expressed as an integer number of contracts. The mechanisms leading to price changes are therefore not expected to vary continuously as prices evolve, but rather to adapt only progressively if prices are seen to rise (or decrease) significantly over a certain time window".

Bouchaud et al. (2001) - studied 437 US stocks, 500 European stocks and 300 Japanese stocks (1990–2000) at the daily frequency using the cross-correlation function, C_{rv} :

$$C_{rv}(\tau) = \frac{Corr(r_{t+\tau}^2, r_t)}{(r_t^2)^2}$$
(1.11)

where r_t is the return at time, t, τ is the time lag. Using the exponential fit:

$$\overline{C_{rv}(\tau)} = -Ae^{(-\tau/T)} \tag{1.12}$$

where A is the amplitude and T is the decay time, they find that: US stocks (A = 1.9 and T = 69 days), European stocks (A = 1.96 and T = 38 days) and Japanese stocks (A = 1.5 and T = 47 days).

The authors remark "this exponential decay should be contrasted with the very slow, power-law like decay of the volatility correlation function, which cannot be characterized by a unique decay time. Therefore, a new time scale seems to be present in financial markets, intermediate between the very high frequency time scale seen on the correlation function of returns (several minutes) and the very low frequency time scales appearing in the volatility correlation function". They only found significant correlations for $\tau \ge 1$ which implies that correlation exists between future volatilities and past price changes or conversely that volatility changes do not convey any useful information on future price changes.

They then developed a simple 'retarded' model where the change of prices are calibrated not on the instantaneous value of the price but on an exponential moving average of the price. The model predicts that $C_{rv}(\tau \rightarrow 0) = -2$ which fits reasonably well for the US and European markets but not for the Japanese markets.

1.5 Volume Hypothesis

Avramov et al. (2006) propose that the Leverage Effect can be fully explained by the interaction between contrarian and herding investors. Their analysis shows that when stock prices fall, herd (uniformed) investors govern the next period volatility. Since they act in the direction of price change they exacerbate the move and cause volatility to increase. However, when the stock price rises, contrarian (informed) investors, govern the next period volatility which causes a reduction in volatility because they trade in the opposite direction to the price change.

They calculate the following regressions for each stock and then group them into quintiles by size:

$$r_t = \sum_{k=1}^{5} a_k D_{k,t} + \sum_{k=1}^{12} b_k r_{t-k} + c \frac{S_t}{N_t} + \varepsilon_t$$
(1.13)

where r_t is the stock return at time, t; D_{kt} are the day-of-the-week dummy variables, S_t is the number of sell transactions at time, t, and N_t is the total number of trades at time, t. $|\varepsilon_t|$ is then the volatility measure for:

$$|\varepsilon_t| = a_0 + a_1 M_t + \sum_{k=1}^{12} b_k |\varepsilon_{t-k}| + cN_t + \left(\delta_0 + \delta_1 \frac{S_{t-1}}{N_{t-1}}\right) \varepsilon_{t-1} + \eta_t$$
(1.14)

where M_t is the Monday dummy variable. 12 lags are used as independent variables to account for the persistence in volatility. The total number of trades, N, is used as an explanatory variable to proxy for the trading volume. The coefficient δ_0 (with $\delta_1 = 0$) has traditionally be been used to study the impact of stock price changes on future volatility. Here they have only considered sell trades but the buy trades can be estimated separately. When $\delta_1 = 0$, the results show that $E[\delta_0] < 0$ for each of the quintiles and for all stocks. This implies that the Leverage Effect is present and strong at the daily frequency. When δ_1 is not restricted to 0, $E[\delta_0] > 0$ and $E[\delta_1] < 0$. A negative δ_1 suggests that the asymmetric volatility effect is time varying with selling activity. The results also indicate that in the presence of a negative unexpected return, selling activity leads to a decrease in the next day volatility. These results are robust to the contemporaneous relationship between selling activity and stock price declines because the variable S_t/N_t ensures that they are orthogonal. The authors conclude "A negative δ_1 combined with a positive δ_0 suggests that the well-documented negative coefficient in the regression of volatility on lagged return, or the asymmetric volatility effect, is entirely attributable to the interaction between selling activity and return".

Having shown the importance of trading activity to the Leverage Effect they then proceed to examine the empirical evidence for their herding/contrarian hypothesis. Sell trades in the presence of positive (negative) unexpected returns were designated contrarian (herding) trades. Formally, the contrarian trades were denoted as, $S_t/N_t * (\varepsilon_t \ge 0)$, where $(\varepsilon_t \ge 0)$ is a dummy variable that is equal to one when the unexpected return is nonnegative and zero otherwise. The herding trades were denoted as, $S_t/N_t * (\varepsilon_t < 0)$, where $(\varepsilon_t < 0)$ is a dummy variable that is equal to one when the unexpected return is negative and zero otherwise. The herding trades were denoted as, $S_t/N_t * (\varepsilon_t < 0)$, where $(\varepsilon_t < 0)$ is a dummy variable that is equal to one when the unexpected return is negative and zero otherwise. The notion is that sell trades in the presence of decreasing prices are designated as herding trades and sell trades in the presence of rising prices are designated as contrarian trades. They assess the impact of contrarian and herding sell trades on volatility using the following specification:

$$|\boldsymbol{\varepsilon}_t| = a_0 + a_1 M_t + \sum_{k=1}^{12} b_k |\boldsymbol{\varepsilon}_{t-k}| + cN_t + \dots$$

$$\left(\delta_{0} + \delta_{1} \frac{S_{t-1}}{N_{t-1}} * (\varepsilon_{t-1} \ge 0) + \delta_{2} \frac{S_{t-1}}{N_{t-1}} * (\varepsilon_{t-1} < 0)\right) \varepsilon_{t-1} + \eta_{t}$$
(1.15)

This specification separates the impact of sell trades by conditioning on positive and

negative unexpected returns. Whilst, it is expected that $\delta_1 + \delta_2 < 0$, the two variables should also each be negative because contrarian sell trades should reduce volatility with $\varepsilon_{t-1} \ge 0$, and herding sell trades should increase volatility with $\varepsilon_{t-1} < 0$. The results show that δ_2 is significantly negative across each quintile for both the number of shares sold and the number of sell trades. This means that when the stock price declines, volatility increases and the increase is attributed to herding sell trades (large and small investors). When selling activity is measured in terms of shares, representing actions of large traders, δ_1 is also significantly negative across all quintiles. This implies that selling activity of large investors leads to a decrease in stock volatility. In addition, since $\delta_2 > \delta_1$, this implies that selling activity has a larger impact on volatility when conditioned on negative unexpected returns. [δ_1 is statistically insignificant when sales are measured in terms of the number of transactions]. The negative coefficient on the herding selling activity and the negative or zero coefficient on the contrarian selling activity combine to give the negative coefficient on the overall selling activity, thereby explaining the asymmetric volatility phenomenon.

They also link herding traders with uniformed traders and contrarian traders with informed traders. They propose that uninformed traders increase volatility whilst informed traders decrease volatility; based upon the works of Hellwig (1980) and Wang (1994). To make the distinction between informed and uninformed traders, they use the theory of Campbell et al. (1993). This posits that sell and buy trades that lead to price reversals can be classed as non-informational trades whilst those that preceded an unchanged stock price are informed trades. The results of this analysis are consistent with that of the contrarian and herding hypothesis. Finally, they analysed the effect of trade size, since Easley and O'Hara (1987) suggests that informed traders are more likely to submit larger orders. Their analysis supports this view and they find that small uninformed trades increase volatility by more than large uniformed trades. Similarly large informed trades reduce volatility by more than small informed trades.

1.5.1 Returns and Volume

The return-volume relation has been studied extensively over the years both at the individual stock and aggregate levels, with much of the early research focused on the contemporaneous relationship. Most of this research focused on well developed markets, usually in the US (Saatcioglu and Starks, 1998). A detailed review of these analyses has been conducted by Karpoff (1987) and Gallant et al. (1992). Whilst Granger and Morgenstern (1963) and Godfrey et al. (1964) were unable to identify any relationships, Ying (1996) discovered that returns and volumes were positively related. This finding has been supported by subsequent research [Epps and Epps (1976), Harris and Gurel (1986), Morgan (1976), Rogalski (1979) and Smirlock and Starks (1988)]. However, some of this research may be misleading because they simply analysed the mean correlation and not the distribution of correlations.

Other research has found that trading volumes can be a useful metric to predict future stock returns. Antoniewicz (1993) found that returns of individual stocks on high volume days are more sustainable than returns on low volume days. Stickel and Verrecchia (1994) found that when earnings announcements are accompanied by higher volume, returns are more sustainable in the following days. Chordia and Swaminathan (2000) found that the returns of stocks with high trading volumes lead those with small volumes. Llorente et al. (2002) found that the returns of stocks of smaller firms show positive auto-correlation and larger stocks show return reversal. Llorente et al. (2002) was able to explain these results by suggesting that hedging (contrarian) trades generate negatively auto-correlated returns. They also went further to relate this to informational asymmetry. Smaller stocks have more information asymmetry and hence more subject to speculative trading.

Recent research has focused on causality and the dynamic relationship between returns and volume. Chuang et al. (2009) state that this is more informative for prediction and risk management. The majority of this research has found a bi-directional relationship between returns and volume. These include Hiemstra and Jones (1994) who used both linear and non-linear Granger causality to examine stocks on the NYSE and Chen et al. (2001) who examined the effect across nine international stock markets. Chuang et al. (2009) noted that examining Granger causality of distribution means (or variances) may lead to misleading results. Hence they examined quantile relationships of the distribution. Using this method they also found two way Granger causality for the NYSE and S&P500 but only volume Granger causes returns on the FTSE100 and not vice versa. Saatcioglu and Starks (1998) in a study of Latin American markets also found that volume causes price changes and not vice versa. However, these results were in conflict with the results of Lee and Rui (2002) who found that volume did not Granger cause returns on the New York, Tokyo or London exchanges.

1.5.2 Volume and Volatility

As reviewed by Karpoff (1987) and Gallant et al. (1992), numerous authors have documented a strong positive relationship between volume and volatility. Crouch (1970a), Crouch (1970b), Karpoff (1987), Wood et al. (1985) and Mulherin and Gerety (1988) found positive correlations between the absolute values of daily price changes and daily volumes. Harris (1983) found similar results for squared values of daily price changes. Whilst Morgan (1976), Westerfield (1977) and Epps and Epps (1976) document that the variance of price changes are positively correlated with volumes. Using evidence of linear and non-linear Granger causality tests, Brooks (1998) found evidence of bi-directional causality with volatility dominating. HoweverJones et al. (1994a)found that the positive volatility-volume relation actually reflects the positive relation between volatility and the number of transactions whilst Chan and Fong (2000a) find that the size of trades and order imbalance play an important role.

There have been models proposed to explain the effect which also apply to the returnvolume relationship. The first are based upon differences in investor opinions and expectations. Wang (1994) developed a rational expectations model which linked trading volume to stock price volatility under asymmetric information, whilst He and Wang (1995) developed a multi-period model with heterogeneous investors and differential information. Harris and Raviv (1993) developed Difference of Opinions theory which assumes that investors are homogenous with respect to their prior beliefs and the new information they receive. The second are information based models. The Mixture of Distributions Hypothesis conjectures that the time series of market returns is drawn from a mixture of conditional distributions with varying degrees of efficiency in generating the expected return. Another model is the sequential information arrival model where Copeland (1976), Copeland (1977) and Jennings et al. (1981) proposed models with asymmetric dissemination of information.

1.5.3 Behavioural Biases and Heuristics

If the Leverage Effect arises due to investor trading behaviour it is also likely that it is influenced by a behavioural process. This proposition is supported by Hens and Steude (2009) and Hasanhodzic and Lo (2011) who state "these results lead us to conclude that rather than being the result of leverage, the inverse relationship between average return and volatility is due to human cognitive perceptions of risk". They suggest that investor behaviour is shaped by their recent experiences which alter their perceptions of risk and hence give rise to changes in demand for risky assets. Hibbert et al. (2008) examined the relationship between stock-market returns and changes in implied volatility (derived from option prices) at both the daily and intraday level. They found that neither the 'Leverage Hypothesis' nor the 'Volatility Feedback Hypothesis' adequately explained the results. They also suggest that their empirical results are consistent with concepts from behavioural finance. However, these behavioural models have only ever been discussed in informal terms.

Gennaioli and Shleifer (2010) have modelled biased perceptions of risk where individuals combine current information with past experiences that are the most representative of the current situation. Such judgments will be biased not only because the representative scenarios that come to mind depend on the situation being evaluated but also because the scenarios that first come to mind tend to be stereotypical ones. In the market context, the first memory that comes to mind for an investor who has experienced significant financial loss is despair. As a result, emotions take hold, prompting the investor to quickly reverse their positions. The view that our recent experiences can have substantial effects on our future behaviour is also backed by Lleras et al. (2009). They show that memories of past experiences affect the kinds of information we pay attention to today. In particular, they compare the effects on the attention system of externally-attributed rewards and penalties to the memory-driven effects that arise when subjects repeatedly perform a task. They find that in both cases the attention system is affected in analogous ways. This leads them to conclude that memories are tainted (positively or negatively) by implicit assessments of our past performance. Other biases of particular relevance are the Representativeness and Affect heuristics (rules of thumb or mental short cuts) and Extrapolation bias. Shefrin (2007) gives an example of Representativeness as when investors judge the risk-return relation for stocks to be negative (based on survey results), since investors view high return and low risk to be representative of a good investment. Hibbert et al. (2008) suggest that this concept can be extended to the market such that larger negative (positive) returns and larger (smaller) risk or volatility are viewed as related characteristics of market behaviour. Finucane et al. (2000a) discusses the 'Affect' heuristic and shows how the labels generated can strongly affect people's decisions. The 'Affect' characteristic is where people form emotional associations with activities, with a positive 'Affect' label being considered good and a negative 'Affect' label being bad. Finally, Extrapolation bias is the extrapolation of past events to form a forecast in combination with those who believe that recent events are representative of the future. Another model that has frequently been used to ascertain how individuals make trading decisions is Prospect Theory. This was developed by Kahneman and Tversky (1979), and posits that investors care about relative changes in wealth rather than absolute wealth and that investors perceive losses and gains differently; losses being more harshly felt than gains (loss-aversion).

1.6 Summary

In this chapter I presented the current research on the Leverage Effect - the negative correlation between an asset's return and its volatility - at the stock level and the evidence for the proposed models. The research, which has largely been focused on western markets, shows that the Leverage Effect is observable at the daily, monthly and quarterly time scales. The main models proposed to explain the Leverage Effect, 'Leverage Hypothesis' (Black, 1976), the 'Volatility Feedback Hypothesis' (numerous authors including Campbell and Hentschel (1992)) and the 'Retarded Volatility Hypothesis' (Bouchaud et al., 2001) all present significant anomalies. The most promising research has been based upon trading volumes. Avramov et al. (2006) propose that, at the stock level, the Leverage Effect arises due to the interaction between contrarian (informed) and herding (uniformed) investors. Herding (uniformed) investors buy stock when the price rises and sell when the price falls which leads to an increase in volatility because they trade in the direction of the price change. It was posited that they govern the volatility dynamics when the stock price falls. Whereas contrarian (informed) investors sell stock when the price rises and buy stock when the price falls which causes a reduction in volatility because they trade in the opposite direction to the price change. It is posited that they govern the volatility dynamics when the stock price rises. This explanation is also consistent with a behavioural based explanation as suggested by Hasanhodzic and Lo (2011) and Hens and Steude (2009), however, this has only ever been discussed informally.

In Chapter 2 I will examine the evidence for trading volumes as an explanation for the Leverage Effect using information theory before analysing a behavioural model in Chapter 3.
Chapter 2

An Information Theoretic Analysis of Stock Returns, Volatility & Trading Volumes

2.1 Introduction

In Chapter 1 I outlined the current research on the Leverage Effect at the stock level. The Leverage Effect refers to the observed negative correlation between an asset's return and its volatility; first documented by Black (1976). Unfortunately, the main economic models, the 'Leverage Hypothesis', the 'Volatility Feedback Hypothesis' and the 'Retarded Volatility Hypothesis', all present significant anomalies. The most promising research (Avramov et al., 2006) asserts that the Leverage Effect can be fully explained by trading volumes. In this chapter I further examine how trading volumes relate to stock returns and volatility for S&P500 stocks. I not only examine the cross-covariance functions, as done by Bouchaud et al. (2001), but also conduct causal inference. Traditionally causal inference has been examined using Granger causality within linear auto-regressive models or non-linear extensions such as those developed by Hiemstra and Jones (1994); this is commonly used among practitioners in finance and economics. Unfortunately, Diks and Panchenko (2005) show that the Hiemstra and Jones (1994) test may not in fact measure

Granger causality and can lead to spurious results. Consequently, I take an alternative approach and use information theory. This is particularly apposite in this context due to its ability to deal with non-linear relationships such as those identified in returns and trading volumes (Hiemstra and Jones, 1994). Furthermore, Barnett et al. (2009) and Hlaváčková-Schindler and Paluš (2007) show that Transfer Entropy (TE), a particular information theoretic measure, is equivalent to Granger causality for a range of distributions. My results highlight the dominant role played by trading volumes in all of the relationships and supports the findings of Avramov et al. (2006) that the Leverage Effect is driven by trading volumes; although I also find that index returns play an important role.

The chapter initially presents a literature review of Granger causality and information theory (Section 2.2) before detailing the data and model calibration (Section 2.3). The results are shown in Section 2.4 with the conclusions given in Section 2.5.

2.2 Literature Review

2.2.1 Granger Causality

Much of the causal inference analysis in the financial and economics literature has relied upon Granger causality. This is based upon the premise that the process X strictly Granger causes another process Y if future values of Y can be better predicted using past values of X rather than only past values of Y. This notion was originally introduced by Wiener (1956) and later formalised in terms of linear auto-regression by Granger (1969). As stated by Barnett et al. (2009), identifying Granger causality is not identical to identifying a physically instantiated causal interaction in a system; this can only be unambiguously identified by perturbing the system. Instead, it is a causal relation in a statistical sense. Hiemstra and Jones (1994) state a condition for instantaneous Granger causality but due to problems in distinguishing between instantaneous causality and instantaneous feedback they only consider strict Granger causality; this is consistent with other research. A major problem with Granger causality is that most real problems, such as the returnvolume relation are non-linear (Hiemstra and Jones (1994) and Chuang et al. (2009)). This led researchers such as Baek and Brock (1992) and Hiemstra and Jones (1994) to develop non-linear extensions to Granger causality. The Hiemstra and Jones (1994) test is now the most commonly used method among practitioners in finance and economics. Unfortunately, Diks and Panchenko (2005) show that this measure may not actually test Granger causality and identify numerous situations in which the test actually fails. An alternative approach is to use a truly non-linear and non-parametric method such as information theory.

Information theory was originally developed to examine properties in signal processing such as data compression by Shannon (1948). However, it is now widely used in the physical sciences for problems such as statistical inference due to its ability to analyse non-linear statistical dependencies. Mutual Information (MI) is a popular measure in the field of information theory. MI gives the mutual reduction in uncertainty of one variable given another. For example, one can calculate the reduction in uncertainty of the daily return at time, t, by knowing the daily volume at time, t. If there is no reduction in uncertainty then the daily returns and volumes are statistically independent. Unfortunately, since this measure is symmetric under the exchange of variables it is only able to determine if two variables are related. However, if one wishes to imply causation one can simply add a time-lag to one variable; this assumes that the causal effect cannot back propagate through time. For example, one can find the reduction in uncertainty in the daily return at time, t + 1, given the daily volume at time, t and vice versa. If there is only a reduction in uncertainty in one direction or one is substantially larger, then one variable must be strongly influencing or causing the changes in the other variable. Alternatively, one can use an asymmetric measure such as TE (Schreiber, 2000), an information theoretic measure of time directed information transfer between jointly dependent processes. Barnett et al. (2009) states that TE is not framed in terms of prediction but in terms of resolution of uncertainty. The TE from Y to X is the degree to which Y disambiguates the future of X beyond the degree to which X already disambiguates its own future. This parallels the notion of Granger causality. In fact Barnett et al. (2009) show that TE is equivalent to Granger causality for Gaussian distributed variables and Hlaváčková-Schindler (2011) extended this to variables distributed as exponential Weinmans, log-normals and certain parametrisations of Generalised Gaussians.

2.2.2 Information Theory

For a continuous random variable, *X*, with probability density, p(x), the differential entropy - the continuous version of the Shannon entropy - is defined as:

$$H(X) = -\int_{X} p(x) \ln p(x) dx \qquad (2.1)$$

this is measured in terms of nats as it uses natural logarithms.

To understand the relationship between two variables one must examine the information transmitted between them. To do this I use MI which for two random variables X and Y is given by:

$$MI(X,Y) = H(X) + H(Y) - H(X,Y)$$
 (2.2)

where H(X, Y) is obtained from the joint distribution p(x, y) of (X, Y).

MI is i) symmetric MI(X,Y) = MI(Y,X), ii) bounded, $0 \le MI(X,Y) \le \min\{H(X),H(Y)\}$, where MI(X,Y) = 0 only if X and Y are independent, and iii) MI(X,Y) = H(Y) only if Y is a function of X. Since MI is symmetric it is necessary to use a time lag (τ) to see how information is transferred through time i.e. $X(t) \rightarrow Y(t+\tau)$ and $Y(t) \rightarrow X(t+\tau)$. Assuming that the effects do not back propagate through time, this method can be used to imply causation between X and Y. The time lagged mutual information, $MI(X;Y)_{t,\tau}$, is given by:

$$MI(X;Y)_{t,\tau} = \int p(x_t, y_{t+\tau}) \ln\left(\frac{p(x_t, y_{t+\tau})}{p(x_t) p(y_{t+\tau})}\right) dxdy$$
(2.3)

Unfortunately, the problem is often complicated by the interaction of additional variables. In order to overcome this problem one can use Partial Mutual Information (PMI), also known as conditional mutual information, which was developed by Frenzel and Pompe (2007). It is based upon the premise of partial correlations which allow one to establish the correlation between two variables, (X, Y), whilst controlling for additional variables, (Z). *PMI* can be seen diagrammatically in Figure 2.1, where the measure of interest is the shaded area, MI(X, Y | Z).



Figure 2.1: Partial Mutual Information (PMI) Diagram

Shows how the Partial Mutual Information, MI(X, Y|Z) (shaded area). It shows the relationship between the entropies, *H*, of variables *X*, *Y* and *Z*.

The PMI is given by

$$MI(X,Y \mid Z) = H(X,Z) + H(Y,Z) - H(Z) - H(X,Y,Z)$$
(2.4)

which corresponds to

$$MI(X,Y \mid Z) = \int p(x,y,z) \ln \frac{p(x,y,z)p(z)}{p(x,z)p(y,z)} dx dy dz$$
(2.5)

where p(x, y, z) is the joint probability of (X, Y, Z). *PMI* is symmetric under the same condition *Z*, MI(X, Y | Z) = MI(Y, X | Z) and $MI(X, Y | Z) \ge 0$; zero is only obtained if *X* and *Y* are independent under condition *Z*. *PMI* can be very useful in controlling for indirect effects. For example, one can examine the relationship between returns and volatility - here represented as *X* and *Y* - whilst controlling for trading volumes - represented by *Z* or one can examine the persistence in the relationship between returns, (*X*), and volumes, (*Y*), whilst controlling for auto-information, (*Z*).

One can also infer causation using a truly asymmetric measure such as TE which gives the time directed information transfer between jointly dependent processes and is a form of Kullback-Leibler divergence (Kullback and Leibler, 1951). Hlaváčková-Schindler and Paluš (2007) show that this is equivalent to a specification of *PMI*. The *TE* from *Y* to *X* is given by:

$$TE(Y \to X)_{t,\tau} = \sum p(x_{t+\tau}, x_t, ..., x_{t-m+2}, y_{t-1}, ..., y_{t-l+1}) \dots$$

$$\ln \frac{p(x_{t+\tau} \mid x_t, \dots, x_{t-m+1}, y_{t-1}, \dots, y_{t-l+1})}{p(x_{t+\tau} \mid x_t, \dots, x_{t-m+1})}$$
(2.6)

where *m* and *l* are the lengths of the vectors *X* and *Y* respectively and τ is the time lag. I have made a slight modification to the standard *TE* formula to condition only on past values of *Y*. This is consistent with Hiemstra and Jones (1994) who identify between strict and instantaneous Granger causality, with the former only conditioning on past values of *Y*.

The *TE* can also be written as:

$$TE(Y \to X)_{\tau} = -H(W_{\tau}, X, Y) + H(W_{\tau}, X) + H(X, Y) - H(X)$$
(2.7)

where $W_{\tau} \equiv x_{t+\tau}$. $TE \ge 0$ and is not symmetric under the exchange of *X* and *Y*. TE = 0 if *X* and *Y* are independent. The Transfer Entropy from *Y* to *X* gives the information about a future observation of *x* obtained from the simultaneous observation of some past of both *x* and *y*, minus the information about the future of *x* obtained from the past of *x* alone.

In a comparative analysis of MI estimators Papana and Kugiumtzis (2009) show that the *k*-nearest neighbour method is the most stable estimator, being less affected by modelspecific parameters. The *k*-nearest neighbour algorithm is a method used to estimate the probability density. In the *k*-nearest neighbour method we grow the volume surrounding the estimation point *x* until it encloses a total of *k* data points. The density then becomes:

$$p(x) = \frac{k}{Nc_d \varepsilon_k^D(x)}$$
(2.8)

where $\varepsilon_k^D(x)$ is the distance between the estimation point *x* and its *k*-th closest neighbour. This can be calculated using a distance measure such as the Euclidian distance. c_d is the volume of the unit sphere in *D* dimensions equal to:

$$c_{d} = \begin{cases} \frac{1}{(D/2)!} \pi^{D/2} & \text{if } D \text{ is even} \\ \frac{1}{(D/2)!} \pi^{(D-1)/2} 2^{D} \left(\frac{D-1}{2}\right)! & \text{if } D \text{ is odd} \end{cases}$$
(2.9)

As outlined by Frenzel and Pompe (2007), this method can be used to calculate the *MI* for two dimensions (X, Y) - which could be returns and volumes or returns and volatility - as follows: For a given number of neighbours, *k*, calculate the maximum distance in the joint distribution, ε_k^2 . Then for each time point, *t*, in the marginal distributions, calculate the number of nearest neighbours, $N_{(t)}$; these are points with the most similar values. This allows one to calculate the harmonic number, $h_{N_{(t)}}$ which is given by:

$$h_{N_{x(t)}} = -\sum_{n=1}^{N} n^{-1}$$
(2.10)

The harmonic numbers may then be used to calculate the MI:

$$MI(X,Y) = \left\langle h_{N_{x(t)}} + h_{N_{y(t)}} \right\rangle - h_{N_{k-1}} - h_{N_{T-1}}$$
(2.11)

where $\langle \rangle$ indicates a time average, *T* is the number of observations and *k* is the prespecified number of neighbours. The *PMI* can also be calculated as:

$$MI(X,Y \mid Z) = \left\langle h_{N_{xz(t)}} + h_{N_{yz(t)}} - h_{N_{z(t)}} \right\rangle - h_{N_{k-1}}$$
(2.12)

It is important to select the appropriate value for k because if k is too small then the estimator will be prone to sampling error but if k is too large then the estimator can be exposed to bias.

2.3 Methodology

2.3.1 The Data

The data is sourced from Bloomberg and covers daily stock returns and volumes for 488 stocks from the S&P500 during the period 1980-2012. The stocks have a daily mean return of 0.0410%, a daily mean standard deviation of 0.0259 and a mean daily volume of 2.2493M shares. The data is not corrected for stocks that have been added/removed from the indices. Since the data represents the constituent stocks as of 2012, the data set is prone to survivorship bias. Stocks that have ceased to trade or have dropped out of the index over this time may display different dynamics/relationships hence this research makes no statement about these stocks. This may be an interesting area of future research.

All statistics and results have been calculated at the individual stock level and then averaged to give a value for the overall index. The significance levels have been estimated by calculating the various measures using surrogate data sets with similar statistical properties but without the inter-relationships; this is consistent with similar research.

2.3.2 Calibrating the Model

First, let me define the variables with which I will conduct the analysis (these are at the daily frequency):

Returns:

$$r_t = \ln[P_t] - \ln[P_{t-1}] \tag{2.13}$$

where P_t is the stock price at time, t.

Normalised Returns:

$$\hat{r}_t = \frac{r_t}{\sigma_r} \tag{2.14}$$

where σ_r is the standard deviation of the returns. Volatility:

$$v_t = r_t^2 \tag{2.15}$$

Normalised Volatility:

$$\hat{v}_t = \frac{v_t}{\sigma_v} \tag{2.16}$$

where σ_v is the standard deviation of the volatility. Normalised Volume:

$$\hat{x}_t = \frac{x_t}{\sigma_x} \tag{2.17}$$

where x_t is the stock volume at time, t, and σ_x is the standard deviation of the volumes. Here I have normalised variables by dividing by the standard deviation of the whole sample of the stock data. The purpose of this is to set each stock to unit variance so that they can easily be compared and to improve the convergence of the estimators. By taking this approach it may bias some techniques because one is influencing historic returns by future volatility. However, the *MI* estimator should in principle be independent of scale factors.

The results are shown with stretched exponential curve fits:

$$f(x) = A \exp\left(-x^{1/T}\right) \tag{2.18}$$

where A is the amplitude, T is the decay rate in days.

In order to calibrate the algorithm it is necessary to ascertain the number of nearest neighbours, k, and the number of observations required to generate stable results. Figure 2.2 shows the $MI(\hat{r}_t, \hat{v}_{t+1})$, $MI(\hat{r}_t, \hat{x}_{t+1})$ and $MI(\hat{x}_t, \hat{v}_{t+1})$. Figure 2.2 (Top) shows that the MI converges as the number of nearest neighbours, k, increases for each of the measures. It indicates that the measures have largely converged for $k \approx 100$. Figure 2.2 (Bottom) shows the MI for varying data lengths from 250 - 2500 days; k = 100. The MI appears stable over the data lengths examined which implies that over these periods

the mean *MI* is stationary; this is a requirement to estimate stable and consistent results. Consequently, in the following analysis I will use k = 100 and exclude stocks with less than 250 observations.



Figure 2.2: Mutual Information Calibration for Stocks at the Daily Frequency

(Top) Shows the mean *MI* for S&P500 stocks as a function of the number of nearest neighbours, *k*. The blue stars with dashed line is the *MI*(\hat{r}_t , \hat{v}_{t+1}), the red diamonds with solid line is the *MI*(\hat{r}_t , \hat{x}_{t+1}) and the green squares with dot-dashed line is *MI*(\hat{x}_t , \hat{v}_{t+1}). These are given with associated one standard errors and exponential curve fits. The results appear to show that the *MI* converges as the number of nearest neighbours increases and that $k \approx 100$ should be sufficient to produce stable results. (Bottom) Shows the mean *MI* for S&P500 stocks across a range of data lengths (days). The blue stars with dashed line is the *MI*(\hat{r}_t , \hat{v}_{t+1}), the red diamonds with solid line is the *MI*(\hat{r}_t , \hat{x}_{t+1}) and the green squares with dot-dashed line is the *MI*(\hat{r}_t , \hat{v}_{t+1}). These are given with associated one standard errors and k = 100. The results indicate that the *MI* does not vary considerably for observation periods between 250 - 2500 days. This suggests that the mean *MI* is stationary over the time periods considered.

2.4 Results

2.4.1 Persistence of Returns, Volatility and Volumes

The first step to understanding the properties of stock returns, volatility and trading volumes is to analyse their auto-mutual information and auto-covariance functions. These functions indicate how the variables persist over time. Figure 2.3 shows the auto-mutual information (Top) and auto-covariance functions (Bottom) for returns, volatility and volumes for S&P500 stocks. It shows that they all exhibit statistically significant autoinformation. It also shows that the auto-information for volumes is significantly larger than for returns or volatility. The auto-covariance functions for volumes and volatility also show that they are highly persistent and statistically significant; with the former being of a larger magnitude. The notable difference is the lack of auto-covariance in returns; this has been documented previously (Cont, 2001). This difference could be an indication that the lower order moments are arbitraged out in the market. This is important as it could lead to spurious results when using linear causality analysis.

- Returns, Volatility and Volumes all display Auto-Information



Figure 2.3: Auto-Mutual Information and Auto-Covariance for Stocks at the Daily Frequency

Shows the auto-mutual information (Top) and normalised auto-covariance (Bottom) functions. They are calculated for returns (blue stars with dashed lines), volatility (green squares with dot-dashed lines) and volumes (red diamonds with solid lines). These are given with associated one standard errors, exponential curve fits and 95% Significance Levels (black dashed lines). The *MI* is statistically significant for returns, volatility and volumes with the *MI* for volumes significantly larger than for returns or volatility. The auto-covariance function exhibits many of the same properties as the auto-mutual information function except that the auto-covariance for returns is not statistically significant at any time horizon.

2.4.2 Returns and Volumes

Figure 2.4 shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for returns and volume. It shows that both the $MI(\hat{r}_t, \hat{x}_{t+\tau})$ and $MI(\hat{x}_t, \hat{r}_{t+\tau})$ are statistically significant and persistent. $MI(\hat{r}_t, \hat{x}_{t+\tau}) > MI(\hat{x}_t, \hat{r}_{t+\tau}) | 1 \le \tau \le 3$ but equal for $\tau > 4$. Hence the *MI* implies a bi-directional information flow between returns and volumes. These results contrast with those of the cross-covariance function which only shows structure for $Cov(r_t, x_{t+\tau})$; this is only statistically significant for $\tau \le 2$. In the past, some authors have found this asymmetry sufficient to imply causation. This would imply that returns cause volumes.

However, when determining causation, the cross-mutual information function can be misleading due to indirect effects such as auto-information. For example, if there is a significant *MI* between returns at time, *t*, and volume at time, *t*, then there could be *MI* between returns at time, *t*, and volume at time, *t* + 1, indirectly due to a significant *MI* between volume at time, *t*, and volume at time, *t* + 1. To correct for auto-information I calculate the partial cross-mutual information function, where I control for the variable $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ where \hat{j} is the normalised returns for $MI(\hat{x}_t, \hat{r}_{t+\tau}|Z)$ and the normalised volumes for $MI(\hat{r}_t, \hat{x}_{t+\tau}|Z)$. Figure 2.5 (Top) shows the partial cross-mutual information function for returns and volume. It shows that only the $MI(\hat{x}_t, \hat{r}_{t+1}|\hat{r}_t)$ is statistically significant at the 95% confidence level. Hence, the *PMI* implies that volumes cause returns and not vice versa. In addition, since neither *MI* is statistically significant for $\tau > 1$, the *PMI* also implies that the persistence is due to auto-information.

Unfortunately, strong and persistent linear correlations, such as those observed in Figure 2.3, can still influence the *PMI* at shorter time horizons. To correct for this I consider the *TE* over longer time horizons as suggested by Schreiber (2000). Figure 2.5 (Bottom) shows the cross-transfer entropy function for returns and volumes. This shows that $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau}) > TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ for all τ . However, they only support the *PMI* for $\tau < 3$ because beyond this point, $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ is statistically significant. Hence the *PMI* may be unduly influenced by the persistence of the linear correlations. The *TE* implies bidirectional (Granger) causality between volumes and returns with volumes dominating.

Persistence in the return-volume relation is driven by auto-information
There is bi-directional (Granger) causality between returns and volumes with volumes dominating



Figure 2.4: Cross-Mutual Information and Cross-Covariance Functions for Stock Returns and Volume at the Daily Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions between returns and volume. The blue squares with solid lines show the $MI(\hat{x}_t, \hat{r}_{t+\tau})$ and $Corr(x_t, r_{t+\tau})$. The green diamonds with dot-dashed lines show the $MI(\hat{r}_t, \hat{x}_{t+\tau})$ and $Corr(r_t, x_{t+\tau})$. These are given with associated one standard errors, stretched exponential curve fits and 95% Significance Levels (dashed black lines). The MI is statistically significant and persistent in both directions with $MI(\hat{r}_t, \hat{x}_{t+\tau}) > MI(\hat{x}_t, \hat{r}_{t+\tau})$ for $1 \le \tau \le 3$ but approximately equal for $\tau > 4$. This indicates bi-directional causality between returns and volumes. However, the cross-covariance function shows that there is only structure for $Corr(r_t, x_{t+\tau})$ and this is only statistically significant for $\tau \le 2$. This could indicate that returns cause volumes and not vice versa.



Figure 2.5: Partial Cross-Mutual Information and Cross-Transfer Entropy Functions for Stock Returns and Volumes at the Daily Frequency

(Top) Shows the partial cross-mutual information function for returns and volume, where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{x}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, \dots, \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised returns. It is given with associated one standard errors, stretched exponential curve fits and a 95% Significance Level (dashed black line). It shows that only the $MI(\hat{x}_t, \hat{r}_{t+1}|\hat{r}_t)$ is statistically significant at the 95% confidence level. Hence, the PMI implies that volumes cause returns and not vice versa. In addition, since neither MI is statistically significant for $\tau > 1$, the *PMI* also implies that the persistence is due to auto-information. (Bottom) Shows the cross-transfer entropy function for returns and volumes. The green diamonds represent $TE(\hat{r}_t \to \hat{x}_{t+\tau})$ and the blue squares represent the $TE(\hat{x}_t \to \hat{r}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \to \hat{x}_{t+\tau})$ and the dashed blue line for $TE(\hat{x}_t \to \hat{r}_{t+\tau})$. The *TE* results show that $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau}) > TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ for all τ . However, they only support the *PMI* results for $\tau < 3$ because beyond this point, $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ is statistically significant. This indicates that the *PMI* may be unduly influenced by the persistence in the linear correlations. The TE implies bi-directional (Granger) causality between volumes and returns with volumes dominating.

2.4.3 Volume and Volatility

Figure 2.6 shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for volumes and volatility. The first thing that you will notice is the similarities between Figures 2.4 and 2.6. This is because the volatility is simply the square of the returns, so the *MI* struggles to differentiate between the two variables. Again, the $MI(\hat{v}_t, \hat{x}_{t+\tau})$ and $MI(\hat{x}_t, \hat{v}_{t+\tau})$ are both statistically significant and persistent. $MI(\hat{v}_t, \hat{x}_{t+\tau}) > MI(\hat{x}_t, \hat{v}_{t+\tau})$ for $1 \le \tau \le 3$ but for $\tau \ge 4$ they are of equal magnitudes. This implies a bi-directional information flow between volumes and volatility. The results of the cross-mutual information are consistent with those of the cross-covariance which show that the $Cov(v_t, x_{t+\tau})$ and $Cov(x_t, v_{t+\tau})$ are both positive and statistically significant for $\tau \le 22$. It is also evident that $Cov(v_t, x_{t+\tau}) > Cov(x_t, v_{t+\tau})$ for $\tau \le 15$.

Figure 2.7 (Top) shows the partial cross-mutual information function for volumes and volatility, where I have controlled for auto-information. It shows that the *MI* is only statistically significant for $MI(\hat{x}_t, \hat{v}_{t+1} | \hat{v}_t)$. Hence the *PMI* implies that volumes cause volatility and not vice versa. Again, since neither *MI* is statistically significant for $\tau > 1$, the *PMI* also implies that the persistence is due to auto-information. The cross-transfer entropy function (Figure 2.7 (Bottom)), indicates that the *PMI* results are again likely the result of persistent linear correlations because the $TE(\hat{v}_t \rightarrow \hat{x}_{t+\tau})$ is statistically significant for $\tau \gtrsim 20$. Therefore the *TE* results indicate bi-directional (Granger) causality between volumes and volatility with volumes dominating.

- Persistence in the volume-volatility relation is driven by auto-information

- There is bi-directional (Granger) causality between volatility and volumes with volumes dominating



Figure 2.6: Cross-Mutual Information and Cross-Covariance Functions for Stock Volatility and Volume at the Daily Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for volumes and volatility. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau})$ and the $Cov(v_t, x_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau})$ and the $Cov(x_t, v_{t+\tau})$. Both graphs are given with associated one standard errors, stretched exponential curve fits and 95% Significance Levels (black dashed lines). The MI is statistically significant and persistent in both directions with $MI(\hat{v}_t, \hat{x}_{t+\tau}) > MI(\hat{x}_t, \hat{v}_{t+\tau})$ for $1 \le \tau \le 3$ but approximately equal for $\tau \ge 4$. This indicates bi-directional causality between volume and volatility. This is consistent with the results of the cross-covariance function which shows that $Cov(v_t, x_{t+\tau}) = Cov(x_t, v_{t+\tau})$ are positive and statistically significant for $\tau \le 22$. It is also evident that $Cov(v_t, x_{t+\tau}) > Cov(x_t, v_{t+\tau})$ for $\tau \le 15$.



Figure 2.7: Partial Cross-Mutual Information and Cross-Transfer Entropy Functions for Stock Volatility and Volumes at the Daily Frequency

(Top) Shows the partial cross-mutual information function for volumes and volatility where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volatility. It is given with associated one standard errors, stretched exponential curve fits and a 95% Significance Level (black dashed line). The MI is only statistically significant for $MI(\hat{x}_t, \hat{v}_{t+1} | \hat{v}_t)$ which indicates that volumes cause volatility and not vice versa. Since neither MI is statistically significant for $\tau > 1$, the PMIalso implies that the persistence is due to auto-information.. (Bottom) Shows the cross-transfer entropy function for volumes and volatility. The green diamonds represent $TE(\hat{v}_t \to \hat{x}_{t+\tau})$ and the blue squares represent $TE(\hat{x}_t \to \hat{v}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels where the green dot-dashed line is for the $TE(\hat{v}_t \to \hat{x}_{t+\tau})$ and the dashed blue line is for the $TE(\hat{x}_t \to \hat{v}_{t+\tau})$. This shows that $TE(\hat{x}_t \to \hat{v}_{t+\tau}) > TE(\hat{v}_t \to \hat{x}_{t+\tau})$ for all τ . However, they only support the results of the PMI for $\tau \leq 20$ because beyond this point, $TE(\hat{v}_t \to \hat{x}_{t+\tau})$ is statistically significant. This indicates that the PMI may be unduly influenced by the persistence of the linear correlations. The TEimplies bi-directional (Granger) causality between volumes and volatility with volumes dominating.

2.4.4 Returns and Volatility

Figure 2.8 shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for returns and volatility. The cross-mutual information function shows that the $MI(\hat{r}_t, \hat{v}_{t+\tau})$ and the $MI(\hat{v}_t, \hat{r}_{t+\tau})$ are both statistically significant and persistent. It also shows that $MI(\hat{r}_t, \hat{v}_{t+\tau}) \approx MI(\hat{v}_t, \hat{r}_{t+\tau})$, as mentioned previously, this is because the MI struggles to differentiate the variables since volatility is simply the square of the returns. This implies a bi-directional information flow between returns and volatility. The Leverage Effect is clearly identifiable in the cross-covariance function by the negative covariance between returns and volatility. The cross-covariance function only shows structure for $Cov(r_t, v_{t+\tau})$; which is statistically significant for $\tau \leq 9$. This led authors, such as Bouchaud et al. (2001), to imply causation from returns to volatility.

Figure 2.9 (Top) shows the partial cross-mutual information function for returns and volatility where I have controlled for auto-information. The *PMI* estimate is overstated because it is not possible to control for the auto-information from the variables when $\tau = 1$ because the volatility is directly calculated from the returns and hence controlling for one indirectly controls for the other. Again the *PMI* is unable to separate returns and volatility but they are now only statistically significant for $\tau \leq 2$ which indicates the persistence is due to auto-information. The *TE* (Figure 2.9 (Bottom)) on the other hand does manage to separate the returns and volatility. It shows that $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau}) > TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ but whilst $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ is only statistically significant for $\tau \approx 10$ days, $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$ is statistically significant for $\tau > 50$ days. This implies bi-directional (Granger) causality between volatility and returns with volatility dominating.

- Persistence in the return-volatility relation is driven by auto-information

- There is bi-directional (Granger) causality between returns and volatility at the daily frequency



Figure 2.8: Cross-Mutual Information and Cross-Covariance Functions for Stock Returns and Volatility at the Daily Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for returns and volatility. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau})$ and the $Cov(r_t, v_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau})$ and the $Cov(v_t, r_{t+\tau})$. Both graphs are given with associated one standard errors, stretched exponential curve fits and 95% Significance Levels (black dashed lines). The cross-mutual information function shows that $MI(\hat{r}_t, \hat{v}_{t+\tau})$ and $MI(\hat{v}_t, \hat{r}_{t+\tau})$ are both statistically significant and persistent. It also shows that $MI(\hat{r}_t, \hat{v}_{t+\tau}) \approx MI(\hat{v}_t, \hat{r}_{t+\tau})$, this is because the MIstruggles to differentiate the variables since the volatility is simply the square of the returns. This implies a bi-directional information flow between returns and volatility. However, the cross-covariance function only shows structure for $Cov(r_t, v_{t+\tau})$; which is statistically significant for $\tau \leq 9$. Some authors have found this sufficient to imply causation from returns to volatility. The negative correlation between returns and volatility is commonly known as the Leverage Effect.



Figure 2.9: Partial Cross-Mutual Information and Cross-Transfer Entropy Functions for Stock Returns and Volatility at the Daily Frequency

(Top) Shows the partial cross-mutual information function for returns and volatility where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 2$ and \hat{j} are the normalised volatilities. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 2$ and \hat{j} are the normalised returns. The squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 2$ and \hat{j} are the normalised returns. It is given with associated one standard errors, stretched exponential curve fits and a 95% Significance Level (black dashed lines). The *PMI* is unable to separate returns and volatility but the *MI* is only statistically significant for $\tau \le 2$ which indicates the persistence is due to auto-information. (Bottom) Shows the cross-transfer entropy function for returns and volatility. The green diamonds represent the $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the blue squares represent $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the dashed blue line for $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. It shows that $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau}) > TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ but whilst $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ is only statistically significant for around $\tau \approx 10$ days, $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$ is statistically significant for $\tau > 50$ days. This implies bi-directional (Granger) causality between volatility and returns with volatility dominating.

In order to decouple the returns and volatility I examined the effects at the weekly frequency where I have defined the volatility as the sum of the squared daily returns over the week. Figure 2.10 shows the cross-mutual information (Top) and normalised crosscovariance (Bottom) functions for returns and volatility at the weekly frequency. Firstly, the cross-covariance function exhibits very similar features at the daily and weekly frequencies, only exhibiting structure for $Cov(r_t, v_{t+\tau})$. Christie (1982), amongst others, has identified a negative correlation between returns and volatility (Leverage Effect) up to the quarterly frequency. Secondly, the *MI* has started to separate the returns and volatility with $MI(\hat{r}_t, \hat{v}_{t+\tau}) > MI(\hat{v}_t, \hat{r}_{t+\tau})$ for $\tau \leq 12$. Since both are statistically significant this implies a bi-directional information flow between returns and volatility. However, the PMI and TE are more interesting. Figure 2.11 (Top) shows the partial cross-mutual information function for returns and volatility at the weekly frequency where I have controlled for auto-information. $MI(\hat{v}_t, \hat{r}_{t+\tau}|Z) > MI(\hat{r}_t, \hat{v}_{t+\tau}|Z)$ for $\tau < 6$, with only $MI(\hat{v}_t, \hat{r}_{t+1}|\hat{r}_t)$ statistically significant. This implies that volatility causes returns and not vice versa. This is supported by the *TE* (Figure 2.11 (Bottom)) but the $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$ is statistically significant for $\tau > 24$ weeks. These results indicate that volatility (Granger) cause returns at the weekly frequency and not vice versa. This could also be true at the daily frequency but in order to prove this one must consider aggregating high frequency data to decouple returns and volatility.

- Volatility (Granger) cause returns at the weekly frequency



Figure 2.10: Cross-Mutual Information and Cross-Covariance Functions for Stock Returns and Volatility at the Weekly Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for returns and volatility at the weekly frequency. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau})$ and the $Cov(r_t, v_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau})$ and the $Cov(v_t, r_{t+\tau})$. Both graphs are given with associated one standard errors, stretched exponential curve fits and 95% Significance Levels (black dashed lines). The cross-mutual information function shows that $MI(\hat{r}_t, \hat{v}_{t+\tau}) > MI(\hat{v}_t, \hat{r}_{t+\tau})$ for $\tau \leq 12$ but they are both statistically significant and persistent. This implies a bi-directional information flow between returns and volatility. The cross-covariance function only shows structure for $Cov(r_t, v_{t+\tau})$ which is statistically significant for $\tau < 5$. This shows that the Leverage Effect is also observable at the weekly frequency.



Figure 2.11: Partial Cross-Mutual Information and Cross-Transfer Entropy Functions for Stock Returns and Volatility at the Weekly Frequency

(Top) Shows the partial cross-mutual information function for returns and volatility at the weekly frequency. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volatilities. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised returns. These are given with associated one standard errors, stretched exponential curve fits and a 95% Significance Level (dashed black line). It shows that the $MI(\hat{v}_t, \hat{r}_{t+\tau}|Z) > MI(\hat{r}_t, \hat{v}_{t+\tau}|Z)$ for $\tau < 6$ with only $MI(\hat{v}_t, \hat{r}_{t+1}|\hat{r}_t)$ statistically significant. This indicates that volatility cause returns at the weekly frequency and not vice versa. (Bottom) Shows the cross-transfer entropy functions for returns and volatility at the weekly frequency. The green diamonds represent the $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the blue squares represent the $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. These are given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the dashed blue line for $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. The *TE* results support those of the *PMI* but the $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$ is now statistically significant for $\tau > 24$ weeks. This indicates that volatility (Granger) cause returns at the weekly frequency.

To examine the impact of trading volumes on the Leverage Effect I examine the *MI* using $MI(\hat{r}_t, \hat{v}_{t+1} | \hat{x}_t, \hat{x}_{t+1})$ (or $MI(\hat{v}_t, \hat{r}_{t+1} | \hat{x}_t, \hat{x}_{t+1})$ since they are equivalent at the daily frequency). I extend the *MI* measure to be an Effective Mutual Information (*EMI*) measure where the *MI* from a secondary 'random' process is subtracted from the *MI* to give the 'significant' information transfer. Table 2.1 shows that trading volumes account for 49.47% of the *EMI* between returns and volatility at the 1 day time lag. However, the *EMI* is still statistically significant. Hence, I also examine if the Leverage Effect is affected by systemic effects by also controlling for the normalised index returns, $\hat{r}_{i,t}$, i.e. $Z = [\hat{x}_t, \hat{x}_{t+1}, \hat{r}_{i,t}, \hat{r}_{i,t+1}]$. Table 2.1 shows that controlling for trading volumes and index returns accounts for 92.63% of the *EMI* between returns and volatility at the 1 day time lag. These results indicate that the relationship between returns and volatility is driven by trading volumes and index returns (evidence of a index level feedback effect).

- Trading volumes are a driver of the return-volatility relation
- Index returns are a driver of the return-volatility relation

Table 2.1: Effective Mutual Information for the Stock Level Leverage Effect Controlling for Trading Volumes and Index Returns

Index	No. of	EMI	EMI	EMI
		$(\hat{r}_t, \hat{v}_{t+1})$	$(\hat{r}_t, \hat{v}_{t+1} \mid \hat{x}_t, \hat{x}_{t+1})$	$(\hat{r}_t, \hat{v}_{t+1} \mid \hat{x}_t, \hat{x}_{t+1}, \hat{r}_{i,t}, \hat{r}_{i,t+1})$
	Stocks			
S&P500	488	0.0095 ± 0.0008	0.0048 ± 0.0005	0.0007 ± 0.0003

Shows the $EMI(\hat{r}_t, \hat{v}_{t+1})$ for S&P500 stocks at the daily frequency. $EMI(\hat{r}_t, \hat{v}_{t+1} | \hat{x}_t, \hat{x}_{t+1})$ shows that trading volumes account for 49.47% of the *EMI* between returns and volatility at the 1 day time lag. Controlling for index feedback effects using index returns, $\hat{r}_{i,t}, \hat{r}_{i,t+1}$. The $EMI(\hat{v}_t, \hat{r}_{t+1} | \hat{x}_t, \hat{x}_{t+1}, \hat{r}_{i,t}, \hat{r}_{i,t+1})$ is now only 0.0007 ± 0.0003 which means that trading volumes and index returns account for 92.63% of the *EMI* between returns and volatility at the 1 day time lag. These results indicate that the relationship between returns and volatility is driven by trading volumes and index feedback effects. [Since the *EMI* cannot differentiate between $EMI(\hat{r}_t, \hat{v}_{t+1})$ and $EMI(\hat{v}_t, \hat{r}_{t+1})$ at the daily frequency, the results presented in this table are equivalent].

2.5 Conclusions

In this chapter I examined the impact of trading volumes on stock returns and volatility for S&P500 stocks. The results highlighted the dominant role played by trading volumes in these relationships. In analysing the Leverage Effect I found that trading volumes accounted for 50% of the *EMI* between returns and volatility. I also found that a further 43% of the *EMI* could be attributed to feedback effects from the index level.

The research also produced a number of stylised facts, from an information theoretic perspective, which may give insights into the functioning of the financial markets:

1) Returns, volatility and volumes all display auto-information.

2) There is bi-directional (Granger) causality between volumes and stock returns but volumes dominate. This supports the findings of Hiemstra and Jones (1994) and Chuang et al. (2009) who also identify bi-directional Granger causality using linear and non-linear Granger methods.

3) At the daily frequency, there is bi-directional (Granger) causality between returns and volatility - with volatility dominating - but at the weekly frequency, volatility (Granger) causes returns and not vice versa.

4) There is bi-directional (Granger) causality between volumes and volatility but volumes dominate. Brooks (1998) also found bi-directional causality using linear and nonlinear Granger causality but he found that volatility dominates.

5) The persistence in the relationships between returns, volatility and volumes are driven by auto-information.

6) Appendix A shows that at the weekly frequency volumes (Granger) cause volatility.

These results are consistent across a range of international markets (Appendix B).

The importance of trading volumes has several implications. Firstly, it implies that we can learn more about the stock market by studying the joint dynamics of stock prices and trading volumes than the univariate dynamics alone; consistent with Gallant et al. (1992). Secondly, we must better understand how investors make trading decisions in order to fully understand the Leverage Effect. Therefore, in the next chapter I investigate the affect of trading behaviour on the relationship between returns and trading volumes.

Chapter 3

Return-Volume Correlation and its Behavioural Origins

3.1 Introduction

In the previous chapter I showed the importance of trading volumes in explaining the stock level Leverage Effect. This means we must understand how we make trading decisions and how these decisions affect stock returns and volatility. In this chapter I use an analytical model to show that the contemporaneous correlation between stock returns and trading volumes is governed by the optimal trading strategy. I also demonstrate that the optimal trading strategy is itself governed by expected stock returns, the standard deviations of returns and investor preferences. I verify these findings empirically using a broad range of developed and emerging market stocks and the trading activity of institutional investors in S&P500 stocks. My findings complement those of Llorente et al. (2002) who show that contrarian (hedging) trades lead to negatively auto-correlated returns.

In order to examine the impact of the optimal trading strategy, I utilise the analytical model of Barberis and Xiong (2009). The model allows me to calculate the optimal trading strategy for investors with Prospect Theory (PT) preferences in a complete market. PT is a model of decision under uncertainty developed by Kahneman and Tversky (1979)

which posits that individuals care about changes in wealth rather than absolute wealth and perceive gains and losses differently. Using this model, I find that return-volume correlation exists when investors have Prospect Theory preferences but not when they have linear preferences. I also find that both negative and positive correlations may exist but they are driven by different trading strategies. Negative return-volume correlation, which is the most prevalent, is generated by a contrarian trading strategy. This is where investors buy as the stock price falls and sell as the stock price rises. The rationale is that since investors are risk-seeking in the loss domain, they increase their stock position when the stock price falls and since they are risk-averse in the gain domain, they reduce their stock position when the stock price rises. On the other hand, positive return-volume correlation is generated by a herding strategy. This is where investors buy as the stock price rises and sell as the stock price falls. This strategy appears optimal when expected returns are high and the standard deviation of returns is low. The rationale is that due to the high expected returns and little downside risk - low standard deviation - they aggressively buy when the stock rises to try to benefit from the rising price.

From the model I draw several predictions that are verified empirically for a broad range of stocks. Return-volume correlation increases as 1) the Sharpe Ratio increases, 2) investors move from contrarian to herding strategies, 3) the curvature of an individual's utility function decreases and 4) the degree of loss-aversion increases. I also re-examine the empirical evidence for the contemporaneous return-volume correlation and extend previous research to a broader range of emerging markets. Whilst the results generally support previous findings of a positive return-volume correlation (summarised in Karpoff (1987) and Gallant et al. (1992)), I also find evidence of negative return-volume correlation; as predicted by the model. The negative correlation appears more prevalent in the US and European markets and at longer time horizons. The reason for this discrepancy is that much of the previous research has focussed on the mean correlation and not considered the distribution of correlations.

The chapter begins by reviewing relevant literature (Section 2.2) and then proceeds to outline the Barberis and Xiong (2009) model and how it relates to return-volume cor-

relation (Section 3.3). The empirical investigation of return-volume correlation and the model predictions is then presented in Section 3.5 with the conclusions in Section 3.6.

3.2 Literature Review

3.2.1 Prospect Theory

Prospect Theory, developed by Kahneman and Tversky (1979), posits that investors care about relative changes in wealth rather than absolute wealth and that investors perceive losses and gains differently; losses being more harshly felt than gains (loss-aversion). Hence it provides a framework to understand how individuals make decisions under uncertainty. It has frequently been used to ascertain how individuals make trading decisions.

Consider the gamble (f, p; g, q), to be read as 'gain f with probability p and g with probability q, independent of other risks' where $f \le 0 \le g$ and p + q = 1.

In the expected utility framework, an investor with utility function $U(\cdot)$ evaluates this risk by computing:

$$U(p, (W+f); q, (W+g)) = pu(W+f) + qu(W+g)$$
(3.1)

where W is his/her current wealth.

By contrast, in the Prospect Theory framework, the investor assigns the gamble the value:

$$V(p,f;q,g) = \phi(p)v(f) + \phi(q)v(g)$$
(3.2)

where $v(\cdot)$ and $\phi(\cdot)$ are known as the value function and the probability weighting function respectively. These functions satisfy v(0) = 0, $\phi(0) = 0$ and $\phi(1) = 1$.

An individual's utility function is then parametrised in the following way:

$$v(f) = \begin{cases} f^{\alpha} & f \ge 0\\ & for \\ -\lambda(-f)^{\beta} & f < 0 \end{cases}$$
(3.3)

for $\alpha, \beta \in (0, 1)$ and $\lambda > 1$.

The utility function is concave in the gain domain which implies that investors are risk-averse with gains because additional utility per unit of gain decreases with the size of the gain. The concavity in the gain domain is governed by α . Conversely the function is convex in the loss domain which implies that investors are risk-seeking with losses because the negative utility per unit of loss decreases with the size of the loss. The convexity in the loss domain is governed by β . It also shows that investors are loss-averse, with losses being more heavily felt than gains ($\lambda > 1$).

Experimentally Kahneman and Tversky (1979) find that $\alpha = \beta = 0.88$ and $\lambda = 2.25$, which implies that the concavity/convexity is mild but there is a strong sense of loss-aversion, with losses felt more than twice as harshly as gains. Interestingly, Abdellaoui et al. (2011) find that finance professionals have utility functions defined by $\alpha = 0.71$ and $\beta = 0.93$. This parametrisation increases the concavity in the gain domain and reduces the convexity in the loss domain. This implies that professionals are more risk-averse in the gain domain but less risk-seeking in the loss domain. Their results for loss-aversion, λ , were mixed but the median value was 1.31, which is significantly less loss-averse than the subjects of Kahneman and Tversky (1979) (Figure 3.1).



Figure 3.1: Prospect Theory Utility Curve

This shows the shape of the utility function for different parametrisations. The black dashed line shows a linear utility function and the blue stars and red squares represent the utility function for students (Kahneman and Tversky, 1979) and professionals (Abdellaoui et al., 2011) respectively. It shows that individuals are risk-averse in the gain domain - function is concave in the gain domain - and risk-seeking in the loss domain - function is close a symmetric, with losses more heavily felt than gains (loss-aversion). It is clear that professionals and students have different utility functions; the former being more risk-averse in the gain domain and more risk-seeking in the loss domain.

3.3 Model (Barberis and Xiong, 2009)

Barberis and Xiong (2009) developed an analytically tractable model of the optimal trading strategy for investors with Prospect Theory preferences. The model assumes a complete market, where the complete set of all possible gambles on future states-of-the-world can be constructed with existing assets, without friction. They built upon the insights of Cox and Huang (1989), who demonstrated that in a complete market, an investor's dynamic optimization problem may be rewritten as a static problem in which he/she directly chooses his/her wealth in the different possible states, at the final date. An optimal trading strategy is then one that generates these optimal wealth allocations; in a complete market such a trading strategy always exists. I will now briefly outline the basis of the model but for an in-depth discussion I refer the reader to the original article.

Consider a portfolio choice setting with dates, t = [0, 1, ..., T], and two assets. A riskfree asset, which earns a gross return of $r_f \ge 1$ in each period, and a risky asset. The price of the risky asset, which may be thought of as an individual stock, evolves as a binomial tree with a transition probability of $\frac{1}{2}$, so there is an equal probability of the stock rising or falling.

To find the optimal strategy, Equation 3.4 maximises the investor's Prospect Theory utility over the period, T (this is taken to be 1 year):

$$V^* = \max_{k \in \{1, \dots, T\}} \left[\left(\sum_{l=1}^k q_{T,l}^{-\frac{\alpha}{1-\alpha}} \pi_{T,l} \right)^{1-\alpha} \left(\sum_{l=k+1}^{T+1} q_{T,l} \pi_{T,l} \right)^{\alpha} - \lambda \sum_{l=k+1}^{T+1} \pi_{T,l} \right]$$
(3.4)

where *T* is the final time, *q* is the price density at node, *l*, and π is the ex-ante probability of reaching node, *l*, in the binomial tree. λ is the loss-aversion parameter and α governs the curvature of the utility function (Equation 3.3). The ex-ante probability, $\pi_{t,j}$, of reaching node, *j*,*t* is given by:

$$\pi_{t,j} = \frac{t! 2^{-t}}{(t-j+1)! (j-1)!}$$
(3.5)

and $q_{t,j}$ is the state price density at node, j,t:

$$q_{t,j} = q_u^{t-j+1} q_d^{j-1}$$
(3.6)

where $q_u = \frac{2(R_f - R_d)}{R_f(R_u - R_d)}$, $q_d = \frac{2(R_u - R_f)}{R_f(R_u - R_d)}$ and $R_{u/d} = \mu^{\frac{1}{T}} \pm \left[\left(\mu^2 + \sigma^2 \right)^{\frac{1}{T}} - \left(\mu^2 \right)^{\frac{1}{T}} \right]^{\frac{1}{2}}$ where R_f is the risk-free return, μ annual gross expected return and σ the standard deviation of the risky asset.

To find the optimal wealth allocations, $W_{t,j}$, and share holdings of the risky asset, $x_{t,j}$, at each node, (t, j), in the tree, one must evaluate the value function (and calculate backwards).

For $V^* > 0$, the optimal wealth allocation, $W_{T,j}$, in node, *j*, at final date, *T*, is given by:

$$W_{T,j} = \begin{cases} W_0 r_f^T \left[1 + q_{T,j}^{-\frac{\alpha}{1-\alpha}} \left(\frac{\sum_{l=k^*+1}^{T+1} q_{T,l} \pi_{T,l}}{\sum_{l=1}^{k^*} q_{T,l}^{-\frac{\alpha}{1-\alpha}} \pi_{T,l}} \right) \right] & if \quad j \le k^* \\ 0 & j > k^* \end{cases}$$
(3.7)

whilst for $V^* \leq 0$

$$W_{T,j} = W_0 r_f^T \ j = 1, ..., T + 1 \tag{3.8}$$

where W_0 is the initial investment and r_f^T is the risk-free return over the period T.

Intermediate wealth allocations (W_t) and share holdings (x_t) are calculated by working backwards from date, *T*:

$$W_{t,j} = \frac{\frac{1}{2}W_{t+1,j}q_{t+1,j} + \frac{1}{2}W_{t+1,j+1}q_{t+1,j+1}}{q_{t,j}} \begin{cases} 0 \le t \le T - 1 \\ for & 1 \le j \le t + 1 \end{cases}$$
(3.9)

$$x_{t,j} = \frac{W_{t+1,j} - W_{t+1,j+1}}{P_0\left(R_u^{t-j+1}R_d^{j-1} - R_u^{t-j+1}R_d^j\right)} \begin{cases} 0 \le t \le T-1 \\ for \\ 1 \le j \le t+1 \end{cases}$$
(3.10)

where P_0 is the initial stock price.

In the case that a strategy offers a non-positive utility, the investor chooses a wealth level of $W_0 r_f^T$ in all final states; as in Equation 3.8. Otherwise, the investor adopts a threshold strategy where he/she allocates a wealth level greater than the risk-free return over the period *T* to the k^* , date, *T*, nodes with the lowest state price densities and a wealth level of zero to the remaining date *T* nodes; as in Equation 3.7.

This model assumes that $\alpha = \beta$ which is consistent with the results of Kahneman and Tversky (1979). For simplicity, the model also ignores the probability weighting function, $\phi(\bullet)$. However, the authors state that this should not affect the analysis because its primary effect is to over-weight low probabilities and this mainly affects assets with highly skewed returns, which stocks are not.

3.3.1 Simulations

In order to examine this model, I simulate N = 100,000 investors trading a stock over a one year period. The investors receive Prospect Theory utility at the end of this period. I use a trading frequency T = 126 which equates to a trading frequency of nearly every other day for a year. T=126 was chosen so that there were a sufficient number of trades by each investor but not so large as to significantly increase the calculation time. [The results do not appear to vary significantly with the trading frequency].

I analyse the degree of return-volume correlation, C_{rx} , by calculating the correlation between the stock return, r, and the change in share holding, δx , over the period t to t + 1, given by:

$$C_{rx} = corr(r_{t,t+1}, \delta x_{t,t+1}) \tag{3.11}$$

where

$$\delta x_{t,t+1} = \ln\left(\frac{x_{t+1}}{x_t}\right) \tag{3.12}$$

In order to evaluate the mechanism driving the volume dynamics, I also calculate the proportion of buy orders, *Z*, that are placed when the stock price rises (falls), as follows:

$$Z_i = \frac{Number \, of \, Buy \, Orders}{Total \, Number \, of \, Orders} \tag{3.13}$$

where *i* indicates the loss or gain domain i.e. when the stock price falls or rises. When $Z_i > 0.5$, investors are more likely to buy stock. Whereas when $Z_i < 0.5$ they are more likely to sell stock.

I now define a metric for the overall trading strategy, S:

$$S = Z_{gain} - Z_{loss} \tag{3.14}$$

When S > 0 investors are using a herding strategy whilst when S < 0 they are using a contrarian strategy.

3.3.2 The Effect of Expected Return and Standard Deviation

For the linear utility function ($\lambda = 1.00$ and $\alpha = 1.00$), there is essentially no correlation between returns and volume for any of the expected returns, μ , or standard deviations, σ , measured (Figure 3.2 (Top)). However, a utility function with Prospect Theory preferences - using the parametrisation ($\lambda = 2.25$ and $\alpha = 0.88$) which is equivalent to that of Kahneman and Tversky (1979) - shows that return-volume correlation does emerge and its degree varies with μ and σ (Figure 3.2 (Bottom)). The fact that return-volume correlation only exists in the case where an investor has Prospect Theory preferences is evidence that return-volume correlation is influenced by ones preferences.

Figure 3.2 (Bottom) shows that negative return-volume correlation (dark regions) exists when the Sharpe Ratio, $\left(\frac{\mu}{\sigma}\right)$, is small and positive return-volume correlation (light regions) when the Sharpe Ratio is large.

Hypothesis 1 : Return-Volume Correlation increases as the Sharpe Ratio $\left(\frac{\mu}{\sigma}\right)$ increases


Figure 3.2: Return-Volume Correlation calculated using the Barberis and Xiong (2009) model for Linear and Prospect Theory Preferences

(Top) Shows the degree of return-volume correlation generated by a linear utility function, for a range of expected returns and standard deviations of returns. There is almost no correlation between returns and volume for the parameter values examined. (Bottom) Shows the degree of return-volume correlation generated by a Prospect Theory utility function, for a range of expected returns and standard deviations. Negative return-volume correlation (blue regions) exists below the diagonal and persist provided σ is sufficient to dominate μ ; otherwise positive return-volume correlation (red and yellow regions) exist. This contrasts with the linear utility function which shows no return-volume correlation. This implies that return-volume correlation is generated by an investor's preferences.

To determine the cause of this return-volume correlation I examine the trading behaviour in the loss and gain domains. Figure 3.3 shows the proportion of buy trades in the loss (Top) and gain (Bottom) domains. The light regions show where investors are more likely to buy (Z > 0.5) and the dark regions where the investors are more likely to sell (Z < 0.5). In the loss domain (Top) investors are more likely to buy when the Sharpe Ratio is small and more likely sell when the Sharpe Ratio is large; effect is reversed in gain domain (Bottom).



Figure 3.3: Trading Behaviour calculated using the Barberis and Xiong (2009) Model for Linear and Prospect Theory Preferences

Shows the proportion of buy orders in the loss (Top) and gain (Bottom) domains, for a range of expected returns and standard deviations. The red and yellow regions show where buy orders dominate and the blue regions where the sell orders dominate. The buy orders dominate in the high volatility and low expected return regimes in the loss domain (Top) but in the low volatility and high expected returns in the gain domain (Bottom). Shows the trading behaviour which results in the return-volume correlation shown in Figure 3.2 (Bottom).

It implies that the return-volume correlation arises due to the imbalance between trading volumes in the loss and gain domains. Negative return-volume correlation arises from contrarian trading strategies where an investor buys stock (or does not sell) when the price falls and sells stock (or does not buy) when the price rises. This arises naturally with Prospect Theory preferences as investors are generally risk-averse in the gain domain and risk-seeking in the loss domain. In the contrarian strategies, buying stock after the price falls dominates. In this case the trading metric, *S*, is negative. On the other hand, positive return-volume correlation arises from herding strategies. This is when an investor buys stock (or does not sell) when the price rises and sells stock (or does not buy) when the price falls. The herding strategy appears to be optimal when the Sharpe Ratio is large. Since the mechanism is driven by buy orders (after the stock price rises), this may be because high expected returns give a good chance of benefiting from a positive return with little downside risk due to the low relative standard deviation. In this case the trading metric, *S*, is positive.

Hypothesis 2: Return-Volume Correlation increases as Investors move from Contrarian to Herding Strategies (i.e. as *S* increases)

3.3.3 The Effect of the Utility Function

First I examine the affect of the curvature of the utility function on return-volume correlation. In the model $\alpha = \beta$, as found by Kahneman and Tversky (1979), so any adjustment to curvature in the gain domain is mirrored by that in the loss domain. The effect of curvature on the utility function is shown in Figure 3.4.



Figure 3.4: Prospect Theory Utility Curve for Varying Curvature

This shows the affect of changing the curvature (α, β) on the shape of the utility function; $(\lambda = 2.25)$. $\alpha = \beta = 0.81$ (blue stars), $\alpha = \beta = 0.88$ (red squares), $\alpha = \beta = 0.95$ (green diamonds) and the dashed black line shows a linear utility function. Decreasing the parameter values leads to an increase in the concavity (convexity) of the utility function over gains (losses).

The affect of curvature on the degree of return-volume correlation is shown in Figure 3.5 (Top). ($\lambda = 1.00$ so that the impact of the convexity adjustment between the gain and loss domains is not asymmetric). At low values of α the investor is strongly risk-averse in the gain domain and strongly risk-seeking in the loss domain. This means that it is always preferable to implement a contrarian trading strategy which gives rise to negative return-volume correlation. The return-volume correlation starts to increase beyond $\alpha > 0.70$; the increase is rapid from $\alpha > 0.80$. This is because the investor becomes less risk-seeking (risk-averse) in the loss (gain) domain which means herding strategies become more attractive. The two effects are in balance around $\alpha = 0.85$. As $\alpha = \beta \rightarrow 1$, the degree of return-volume correlation tends back towards zero which is consistent with the previous

finding for a linear utility function. We can infer that an increase in α increases the impact of high Sharpe Ratios, which drive herding strategies, and reduces the prevalence of contrarian strategies which are associated with negative return-volume correlation.

Hypothesis 3: Return-Volume Correlation increases as the Curvature of an Individual's Utility Function decreases (i.e. as α increases)

Next I examine the affect of loss-aversion, λ , on return-volume correlation. The degree of return-volume correlation increases approximately linearly with the degree of lossaversion, λ , (Figure 3.5 (Bottom)). This is because increasing the degree of loss-aversion, λ , serves to reduce the impact of convexity in the loss domain and hence increases the degree of return-volume correlation.

Hypothesis 4: Return-Volume Correlation increases as the degree of Loss-Aversion (λ) increases



Figure 3.5: Effect of Varying Curvature and Loss Aversion on Return-Volume Correlation in the Barberis and Xiong (2009) Model

(Top) This shows the effect of different curvatures on return-volume correlation. The degree of return-volume correlation appears to be relatively unaffected by the curvature until ($\alpha = \beta > 0.82$) at which point it increases dramatically. As $\alpha = \beta \rightarrow 1$ the return-volume correlation disappears as shown in Figure 3.2 for a linear utility function .[T = 126, $\sigma = 0.3$, $\mu = 1.03$, $\lambda = 2.25$, N = 100,000]. (Bottom) This shows the effect of loss-aversion (λ) on return-volume correlation. The degree of return-volume correlation appears to be positively correlated with λ and observes a linear relationship. [T = 126, $\sigma = 0.3$, $\mu = 1.03$, $\alpha = \beta = 0.88$, N = 100,000]

3.4 Methodology

3.4.1 The Data

In order to conduct the analysis I have sourced daily price and volume data from Bloomberg for stocks from a range of global indices. These stocks have then been separated into developed and emerging markets. The developed markets consist of Australia (ASX), Canada (TSX), Europe (ESTX), France (CAC), Germany (DAX), Holland (AEX), Hong Kong (HSI), Italy (MIB), Japan (NIKKEI), South Korea (KOSPI), Spain (IBEX), Sweden (OMX), Switzerland (SMI), UK (FTSE) and the US (DJIA, NASDAQ, S&P500). The emerging markets consist of Argentina (MERVAL), Brazil (BOVESPA), China (SHANG-HAI, SHENZEN, HSCEI), India (SENSEX) and Mexico (MEXBOL). The data covers the period 2003-2013. I also have data for the weekly aggregate buy/sell orders executed by institutional shareholders in S&P500 stocks over the period 2010-2014 and the implied volatility data for the S&P500, individual S&P500 stocks and the VIX for the period 2005-2014.

The data is not corrected for stocks that have been added/removed from the indices. Since the data represents the stocks from different indices as of 2012, the data set is prone to survivorship bias. Stocks that have ceased to trade or have dropped out of the index over this time may display different dynamics/relationships hence this research makes no statement about these stocks. This may be an interesting area of future research. The only filter that has been used, is to remove stocks that have an insufficient number of observations for the required observation frequency.

3.5 Results

3.5.1 Return-Volume Correlation

Firstly, it is often useful to view the data. Figure 3.6 shows the density plots for the returnvolume correlation grouped by developed and emerging market stocks. It is separated into annual, quarterly and monthly timescales. From these plots several features are apparent. Firstly, both positive and negative return-volume asymmetry exist at all scales. Hence it can be misleading to simply quote the mean return-volume correlation as has been done in previous research (see Section 2.2). Secondly, the return-volume correlation is more positively skewed in emerging market stocks than developed market stocks and thirdly, the range of return-volume correlation increases with observation frequency. I will now proceed to examine the model predictions.



Figure 3.6: Density Plots of Return-Volume Correlation for Developed and Emerging Market Stocks

This shows the density plots of the of the return-volume correlation for developed (left) and emerging (right) market stocks. They have been separated by observation scale, annual (Top), quarterly (Middle) and monthly (Bottom). From these plots several features are readily apparent. Firstly, both positive and negative return-volume asymmetry exist at all scales. Secondly, the return-volume correlation is more positively skewed in emerging market stocks than developed market stocks. Thirdly, the range of return-volume correlation increases with observation frequency.

3.5.2 Model Predictions

Hypothesis 1: Return-Volume Correlation increases as the Sharpe Ratio $\left(\frac{\mu}{\sigma}\right)$ increases

To examine the affect of the Sharpe Ratio on return-volume correlation I analyse the linear regression:

$$C_{rx,i} = a_0 + a_1 \frac{r_i}{\sigma_i} + \varepsilon \tag{3.15}$$

where $C_{rx,i}$ is the degree of return-volume correlation for a particular index, *i*, and is calculated as the correlation between returns at time, *t*, and volumes at time, *t*. *r* and σ are the mean daily returns and the standard deviations of daily returns respectively, over the given period.

Table 3.1 shows the regression coefficients for stocks across a broad range of international equity markets. Panels A, B and C are for annual, quarterly and monthly observations respectively. In each index, return-volume correlation is correlated positively with the Sharpe Ratio and this correlation is highly statistically significant; t-values are shown in brackets. This is entirely consistent with the model predictions.

Table 3.1: Regression for Return-Volume Correlation andSharpe Ratio for Developed Market Indices

Panel A: Annual Data

		Si	ummary Stat	istics			Regression 1	Results
Index	No.	No.	Daily $\bar{\mu}$	Daily	Daily \bar{x}	$\bar{C_{rx}}$	Constant	Sharpe
	of	of	(%)	$\bar{\sigma}$	(M)			Ratio
	Stocks	Obs.						
ASX	200	1656	0.032	0.027	0.611	0.036	0.021 (6.0)	0.864 (16.7)
TSX	246	1956	0.038	0.024	3.442	0.035	0.018 (5.8)	0.762 (16.2)
ESTX	50	495	0.012	0.020	1.386	-0.015	-0.029(-5.4) 0.952(10.4)
CAC	40	388	0.010	0.021	6.413	-0.013	-0.029(-4.5) 1.192(10.6)
DAX	30	297	0.030	0.021	1.386	-0.010	-0.039(-4.9)) 1.254(9.9)
AEX	24	221	0.010	0.021	1.048	-0.026	-0.043(-4.4)) 1.267(7.6)
HSI	49	394	0.056	0.024	5.212	0.064	0.034 (4.4)	0.959 (9.4)
MIB	40	337	-0.008	0.020	7.776	0.050	0.046(7.1)	0.674(6.2)
NIKKEI	225	1945	-0.005	0.023	1.585	0.101	0.095 (35.5)	1.500 (32.5)
IBEX	35	289	0.000	0.020	50.733	0.004	-0.002(-0.2)) 0.347 (3.8)
OMX	30	262	0.035	0.022	10.970	-0.006	-0.030(-3.3) 1.092(7.8)
KOSPI	760	5856	0.026	0.031	0.033	0.132	0.123 (68.0)	0.908 (33.1)
SMI	20	166	0.016	0.019	0.797	-0.022	-0.044(-3.8)) 1.337(6.9)
FTSE	99	923	0.032	0.020	10.483	0.006	-0.018(-3.9)) 1.018(11.9)
DJIA	30	270	0.016	0.018	3.472	-0.043	-0.066(-7.1)) 1.296(7.7)
NASDAQ	97	793	0.060	0.025	7.208	-0.009	-0.059(-9.5)) 1.913 (18.6)
S&P500	488	4145	0.030	0.022	1.137	-0.022	-0.050(-20.0)	0)1.289(30.1)

		Su	immary Stat	istics			Regression	Results
Index	No.	No.	Daily $\bar{\mu}$	Daily	Daily \bar{x}	$\bar{C_{rx}}$	Constant	Sharpe
	of	of	(%)	$\bar{\sigma}$	(M)			Ratio
	Stocks	Obs.						
ASX	200	6,868	0.029	0.026	0.447	0.046	0.030(12.7)	0.774(39.3)
TSX	248	8,445	0.035	0.023	2.794	0.050	0.030(14.0)	0.701(40.5)
ESTX	50	1,987	0.011	0.019	1.180	-0.001	-0.023(-5.4) 1.001(26.7)
CAC	40	1,561	0.010	0.021	6.690	0.003	-0.020(-4.1) 1.140(26.0)
DAX	30	1,190	0.031	0.020	1.180	0.011	-0.021(-3.4) 1.002(21.8)
AEX	25	891	0.010	0.020	1.198	0.002	-0.022(-3.1) 1.141(18.7)
HSI	49	1,701	0.054	0.022	5.059	0.073	0.036(7.4)	1.163(30.3)
MIB	40	1,370	-0.006	0.020	5.200	0.067	0.055(10.3)	0.900(21.8)
NIKKEI	225	8,220	-0.004	0.022	1.824	0.103	0.096(47.9)	1.109(66.2)
IBEX	35	1,166	0.001	0.019	30.793	0.012	-0.000(-0.0) 0.565(13.8)
OMX	30	1,134	0.037	0.021	9.380	0.017	-0.021(-3.1) 1.217(21.2)
KOSPI	771	25,243	0.025	0.030	0.002	0.140	0.134(100.4)	0.921(84.8)
SMI	20	720	0.020	0.018	0.544	0.016	-0.021(-2.7) 1.244(19.7)
FTSE	101	3,730	0.032	0.020	8.282	0.029	0.001(0.2)	0.858(30.0)
DJIA	30	1,170	0.021	0.017	2.283	-0.008	-0.041(-6.7) 1.151(22.9)
NASDAQ	99	3,438	0.059	0.024	4.467	0.023	-0.031(-7.1) 1.559(43.1)
S&P500	499	17,988	0.033	0.021	0.938	0.010	-0.029(-16.	7)1.166(82.6)

Panel B: Quarterly Data

		Su	mmary Stat	istics			Regression F	Results
Index	No.	No.	Daily $\bar{\mu}$	Daily	Daily \bar{x}	$\bar{C_{rx}}$	Constant	Sharpe
	of	of	(%)	$\bar{\sigma}$	(M)			Ratio
	Stocks	Obs.						
ASX	200	20,736	0.030	0.025	0.888	0.052	0.038(20.3)	0.635(69.1)
TSX	248	25,744	0.037	0.022	3.119	0.060	0.042(24.7)	0.564(70.2)
ESTX	50	6,001	0.013	0.019	0.908	0.016	-0.008(-2.2)	0.831(48.5)
CAC	40	4,724	0.012	0.020	3.473	0.025	0.001(0.1)	0.949(47.4)
DAX	30	3,601	0.032	0.020	0.908	0.027	-0.004(-0.7)	0.812(38.4)
AEX	25	2,699	0.011	0.020	0.819	0.021	-0.005(-1.0)	0.907(34.7)
HSI	49	5,122	0.054	0.022	4.569	0.080	0.047(12.4)	1.027(56.8)
MIB	40	4,118	-0.006	0.019	4.709	0.083	0.068(15.6)	0.742(39.1)
NIKKEI	225	25,082	-0.007	0.022	1.439	0.110	0.101(59.7)	0.800(106.0)
IBEX	35	3,514	0.003	0.019	17.412	0.030	0.015(3.4)	0.564(27.4)
OMX	30	3,459	0.037	0.021	5.240	0.026	-0.003(-0.5)	0.800(31.7)
KOSPI	774	76,369	0.026	0.029	0.003	0.132	0.129(120.4)	0.822(166.3)
SMI	20	2,198	0.021	0.018	0.364	0.029	-0.002(-0.3)	0.930(33.7)
FTSE	101	11,216	0.032	0.019	11.539	0.039	0.015(5.6)	0.656(50.9)
DJIA	30	3,570	0.021	0.016	1.660	0.006	-0.022(-4.7)	0.867(40.2)
NASDAQ	100	10,496	0.059	0.023	3.736	0.039	0.001(0.3)	1.003(70.6)
S&P500	500	54,881	0.034	0.020	0.444	0.023	-0.007(-5.7)	0.840(143.3)

Shows the linear regression: $C_{rx} = a_0 + a_1 \frac{r}{\sigma} + \varepsilon$ for a range of developed market indices (2002-2012), where C_{rx} is the correlation between returns at time, *t*, and volumes at time, *t*. *r* and σ are the returns and the standard deviations of returns respectively, over the given period. t-values are shown in brackets. Panels A, B and C reflect the annual, quarterly and monthly observations respectively. C_{rx} is correlated positively with the Sharpe Ratio for all markets at each observation frequency and is highly statistically significant. Results are consistent with the model predictions. The summary statistics are shown as averages denoted by the symbol.

Table 3.2: Regression for Return-Volume Correlation andSharpe Ratio for Emerging Market Indices

Panel A	A: Annual	Data
---------	-----------	------

		Sı	Regression Results				
Index	No.	No.	Daily $\bar{\mu}$	Daily	Daily \bar{x}	$\bar{C_{rx}}$	Constant Sharpe
	of	of	(%)	$\bar{\sigma}$	(M)		Ratio
	Stocks	Obs.					
MERVAL	12	94	0.072	0.028	0.039	0.093	0.073 (4.6) 0.749 (3.9)
BOVESPA	69	497	0.059	0.028	18.532	0.048	0.033(6.3) 0.585(8.1)
HSCEI	38	240	0.068	0.030	6.681	0.094	0.075 (9.5) 0.681 (6.5)
SHANGHA	I 957	7585	0.012	0.030	85.535	0.214	0.214(190.1) 0.193(12.3)
SHENZEN	1378	6553	-0.003	0.031	0.710	0.202	0.203 (172.3) 0.290 (17.3)
SENSEX	30	253	0.098	0.024	0.142	0.073	0.028(3.1) 1.021(9.1)
MEXBOL	34	266	0.080	0.021	17.982	0.067	0.050(5.6) 0.376(3.5)

Panel B: Quarterly Data

		Su	mmary Stati	stics			Regression Results	
Index	No.	No.	Daily $\bar{\mu}$	Daily	Daily \bar{x}	$\bar{C_{rx}}$	Constant Sharpe	
	of	of	(%)	$\bar{\sigma}$	(M)		Ratio	
	Stocks	Obs.						
MERVAL	12	402	0.063	0.027	0.084	0.125	0.101(10.5) 0.856(12.3)	
BOVESPA	69	2,126	0.056	0.026	21.023	0.064	0.044(11.2) 0.603(19.8)	
HSCEI	40	1,040	0.058	0.028	7.455	0.103	0.078(13.5) 1.006(22.1)	
SHANGHA	I 985	31,530	0.006	0.029	50.404	0.219	0.222(241.4) 0.418(61.2)	
SHENZEN	1,520	28,372	-0.010	0.030	0.572	0.204	0.209(217.6) 0.459(62.6)	
SENSEX	30	1,097	0.090	0.023	0.150	0.092	0.044(7.0) 0.999(21.7)	
MEXBOL	35	1,099	0.080	0.020	16.027	0.092	0.065(10.3) 0.509(11.3)	

		Regression Results						
Index	No.	No.	Daily $\bar{\mu}$	Daily	Daily \bar{x}	$\bar{C_{rx}}$	Constant	Sharpe
	of	of	(%)	$\bar{\sigma}$	(M)			Ratio
	Stocks	Obs.						
MERVAL	12	1,213	0.062	0.026	0.036	0.133	0.112(14.2)	0.690(20.0)
BOVESPA	69	6,419	0.057	0.026	15.517	0.072	0.052(15.9)	0.587(39.6)
HSCEI	40	3,155	0.057	0.028	7.702	0.097	0.076(15.6)	0.960(42.9)
SHANGHA	I 992	95,606	0.005	0.028	42.533	0.210	0.210(229.6) 0.031(33.0)
SHENZEN	1,549	86,629	-0.011	0.029	1.091	0.192	0.195(229.0)) 0.598(163.0)
SENSEX	30	3,318	0.089	0.022	0.079	0.118	0.076(15.4)	0.794(36.7)
MEXBOL	35	3,322	0.079	0.020	16.951	0.100	0.073(14.8)	0.477(22.7)

Panel C: Monthly Data

Shows the linear regression: $C_{rx} = a_0 + a_1 \frac{r}{\sigma} + \varepsilon$, for a range of emerging market indices (2002-2012), where C_{rx} is the correlation between returns at time, *t*, and volumes at time, *t*. *r* and σ are the returns and the standard deviations of returns respectively, over the given period. t-values are shown in brackets. Panels A, B and C reflect the annual, quarterly and monthly observations respectively. C_{rx} is correlated positively with the Sharpe Ratio for all markets at each observation frequency and is highly statistically significant. Results are consistent with the model predictions. The summary statistics are shown as averages denoted by the symbol.

In this analysis I have time-series data for a range of stocks. Two-dimensional data sets of this type are often referred to as panel data. Since effects may not be independent across time or across stocks it is necessary to control for these effects. To do this one may use panel regressions where the standard errors are clustered by index and time (Thompson, 2011). These are shown in Table 3.3 where the t-values are given in brackets for each of the regressions. The results shows that the Sharpe Ratio remains positive at each timescale and is highly statistically significant for all the regressions, except at the monthly timescale, where the standard errors are clustered by date and index and only by index. These findings are also consistent with the model predictions.

Table 3.3: Panel Regression for Return-Volume Correlation

and Sharpe Ratio

	Fixed		Annual	Quarterly	Monthly	
	Effect	ts?				
	Time	Index	Constant Sharpe	Constant Sharpe	Constant Sharpe	
			Ratio	Ratio	Ratio	
Coefficient	Y	Y	-0.043 0.707	-0.030 0.779	0.106 0.131	
Estimate						
t-stat (Std. Errors	Y	Y	(-5.0) (5.7)	(-5.2) (8.1)	(11.9) (1.3)	
clustered by Index						
and Time)						
t-stat (Std. Errors	Y	Y	[-9.6] [5.8]	[-3.56] [8.4]	[6.3] [1.3]	
clustered by Index)						
t-stat (Std. Errors	Y	Y	$\{-2.9\}$ $\{12.7\}$	{-3.5} {21.5}	{14.2} {2.1}	
clustered by Time)						
$r^{2}(\%)$	Y	Y	40.89	30.57	12.17	

Shows the linear regression between return-volume correlation and the Sharpe Ratio for combined developed and emerging market stocks (2002-2012), at the annual, quarterly and monthly timescales. It shows the t-values for the regressions when the standard errors are clustered by date and by index. All regressions control for both time and index fixed effects. The results show that the return-volume correlation and the Sharpe Ratio are positively correlated at each timescale. The correlation is statistically significant for all regressions except at the monthly scale when the standard errors are clustered by index and date effects and by only index. Results are consistent with the model predictions.

Hypothesis 2: Return-Volume Correlation increases as Investors move from Contrarian to Herding Strategies (i.e. as *S* increases)

To analyse the trading behaviour associated return-volume correlation I can examine the aggregate buy/sell orders executed by institutional shareholders. I estimate the aggregate trading strategy used by institutional investors on a weekly basis over a given period and estimate the following linear regression:

$$C_{rx} = a_0 + a_1 S + \varepsilon \tag{3.16}$$

where C_{rx} is the degree of return-volume correlation, calculated as the correlation between returns at time, *t*, and volumes at time, *t*. *S* is the trading strategy metric which is positive for a herding strategy and negative for a contrarian strategy; it is estimated from Equation 3.14.

Table 3.4 shows the regression coefficients for the S&P500 with annual, quarterly and monthly observations. The t-values are shown in brackets. This shows that *S* increases with C_{rx} at each timescale and that these correlations are highly statistically significant. This is consistent with the model predictions.

			Summar	y Statistic	S		Regre	ession Resu	lts
Data	No.	No.	Daily	Daily	Daily \bar{x}	$\bar{C_{rx}}$	Constant	S	$r^{2}(\%)$
	of	of	μ <u>¯</u> (%)	$\bar{\sigma}$	(M)				
	Stocks	Obs.							
Annual	494	4,130	0.030	0.021	1.924	-0.020	-0.059	0.179	2.10
							(-17.6)	(6.4)	
Quarterly	499	16,627	0.031	0.020	1.920	0.003	-0.034	0.172	3.39
							(-12.4)	(16.3)	
Monthly	499	50,788	0.032	0.019	1.906	0.009	-0.034	0.181	5.10
							(-16.2)	(35.1)	

Table 3.4: Regression for Return-Volume Correlation andTrading Strategy (S) for S&P500 Institutional Trading

Shows the linear regression: $C_{rx} = a_0 + a_1S + \varepsilon$, for S&P500 stocks (2010-2014), where C_{rx} is the correlation between returns at time, *t*, and volumes at time, *t*. *S* is the trading metric which is positive for a herding strategy and negative for a contrarian strategy; it is estimated from Equation 3.14. The t-values are shown in brackets. At each time horizon, the return-volume correlation is positively correlated with the trading strategy metric, *S*. The results are statistically significant at both the quarterly and monthly timescales. Results are consistent with the model predictions. The summary statistics are shown as averages denoted by the symbol.

Hypotheses 3 and 4: Return-Volume Correlation increases as the Curvature of an Individual's Utility Function decreases (i.e. as α increases) and as the degree of Loss-Aversion (λ) increases

Ascertaining the impact of the utility function is significantly more challenging since it is not obvious how to extrapolate the aggregate utility function active in empirical stock data. One possibility is to segment the stocks based upon the proportion of active professional investors. This is because it has been shown that professionals and non-professionals (students) have different shaped utility functions and degrees of lossaversion; the former being more risk-averse in the gain domain and more risk-seeking in the loss domain due to lower loss-aversion (Abdellaoui et al., 2011). Hence, the degree of return-volume correlation should be smaller in markets with higher proportions of professional investors. This should be evident by comparing developed and emerging market stocks since professional investors are more active in the former. Table 3.5 shows the mean return-volume correlation on a per market basis grouped into developed and emerging markets. It is clear from these results that the developed market indices have lower degrees of return-volume correlation. However, as stated previously, it can be misleading to simply look at mean values. Hence, Figure 3.6 also shows us that the distribution of the return-volume correlations is more positively skewed for emerging market stocks. These results are consistent with the proposition that there are more active professional investors in developed markets and the model predictions.

The fact that return-volume correlation decreases with the proportion of professional investors is consistent with the findings of Talpsepp and Reiger (2010). They studied the Leverage Effect in 49 countries and found that it is more negative in more developed countries. However, they posit that this is caused by a larger number of private investors which they estimate by analysing the number of households owning stock. Unfortunately, they fail to account for the fact that these private investors represent a far smaller proportion of the daily trading activity in developed markets.

Table 3.5: Return-Volume Correlation Grouped into Devel-

oped and Emerging Markets

Annual	Data	Quarterly	y Data	Monthly Data		
Developed	Emerging	Developed	Emerging	Developed	Emerging	
KOSPI[0.132]	SHANGHAI[0.21	4]KOSPI[0.140]	SHANGHAI[0.219	9]KOSPI[0.132]	SHANGHAI[0.210]	
NIKKEI[0.10]	I]SHENZEN[0.202]	NIKKEI[0.103	3]SHENZEN[0.204]	NIKKEI[0.110]SHENZEN[0.192]	
HSI[0.064]	HSCEI[0.094]	HSI[0.073]	MERVAL[0.125]	MIB[0.083]	MERVAL[0.133]	
MIB[0.050]	MERVAL[0.093]	MIB[0.067]	HSCEI[0.103]	HSI[0.080]	SENSEX[0.118]	
ASX[0.036]	SENSEX[0.073]	TSX[0.050]	MEXBOL[0.092]	TSX[0.060]	MEXBOL[0.100]	
TSX[0.035]	MEXBOL[0.067]	ASX[0.046]	SENSEX[0.092]	ASX[0.052]	HSCEI[0.097]	
FTSE[0.006]	BOVESPA[0.048]	FTSE[0.029]	BOVESPA[0.064]	FTSE[0.039]	BOVESPA[0.072]	
IBEX[0.004]		NASDAQ[0.02	23]	NASDAQ[0.03	39]	
OMX[-0.006]	OMX[0.017]		IBEX[0.030]		
NASDAQ[-0.	.009]	SMI[0.016]		SMI[0.029]		
DAX[-0.010]		IBEX[0.012]		DAX[0.027]		
CAC[-0.013]		DAX[0.011]		CAC [0.026]		
EUROSTOXX	K[-0.015]	S&P500[0.010)]	OMX[0.026]		
S&P500[-0.0	22]	CAC[0.003]		S&P500[0.023]	
SMI[-0.022]		AEX[0.002]		AEX[0.021]		
AEX[-0.026]		EUROSTOXX	X[-0.001]	EUROSTOXX	[0.016]	
DJIA[-0.043]		DJIA[-0.008]		DJIA[0.006]		

Shows mean return-volume correlation grouped into developed and emerging markets. It is lower in the developed markets which is consistent with the predictions of the model given that developed markets have a larger proportion of professional investors who are more prone to contrarian trading strategies. This was inferred from Abdellaoui et al. (2011) who showed professional investors to be more risk-averse in the gain domain and more risk-seeking in the loss domain, when considering lower degrees of loss-aversion.

Hypothesis 4: Return-Volume Correlation increases as the degree of Loss-Aversion (λ) increases

I examine loss-aversion, λ , indirectly through the use a proxy variable such as the VIX index; this is commonly used a proxy for the degree of 'fear' in the market. I also use the at-the-money forward (ATMF) implied volatility for the S&P500 and the individual stock. All of these proxy measures should act in a similar fashion to loss-aversion, λ , and hence be positively correlated with the degree of return-volume correlation, C_{rx} . I test this using the following linear regression:

$$C_{rx} = a_0 + a_1 \frac{r}{\sigma} + a_2 L + \varepsilon \tag{3.17}$$

where C_{rx} is the degree of return-volume correlation and is calculated by the correlation between returns at time, t, and volumes at time, t. r and σ are the returns and the standard deviations of returns respectively, over the given period. L is the proxy measure for loss-aversion which can be the VIX, the ATMF IV of the individual stock or the ATMF IV of the S&P500.

Table 3.6 shows the regression coefficients for S&P500 stocks for each proxy measure; the t-values are given in brackets. It shows that all of the proxy measures, L, are positively correlated with C_{rx} and are statistically significant. This is also consistent with the model predictions.

Table 3.6: Regression for Return-Volume Correlation and

Loss-Aversion (λ) for S&P500 Stocks

Summary			Regressio				
	Statistic	Statistics					
Loss-Aversion	No. of	Obs.	-	Constant	Sharpe	L	r^2
Proxies	Stocks				Ratio		(%)
Individual Stock	494	4,055	-	-0.120(-19.8)	1.100(28.0)	0.002(13.5)	17.04
ATMF IV							
S&P500 ATMF IV	494	4,055		-0.090(-11.5)	1.045(25.7)	0.002(6.1)	14.10
VIX	494	4,055		-0.089(-12.0)	1.053(25.8)	0.002(6.3)	14.73

Panel A: Annual Data Regression Results

Panel B: Quarterly Data Regression Results

	Summary		Regressio	Regression Results		
	Statistic	s				
Loss-Aversion	No. of	Obs.	Constant	Sharpe	L	r^2
Proxies	Stocks			Ratio		(%)
Individual Stock	494	16,388	-0.106(-26.3)	1.118(76.5)	0.002(19.7)	26.46
ATMF IV						
S&P500 ATMF IV	494	16,388	-0.079(-16.8)	1.098(74.2)	0.002(10.3)	25.20
VIX	494	16,388	-0.075(-17.5)	1.102(74.1)	0.002(10.5)	25.22

Panel C: Monthly	Data Regress	sion Results
------------------	--------------	--------------

	Summary S	Statistics	Regressi	on Results		
Loss-Aversion	No. of	Obs.	Constant	Sharpe	L	r^2
Proxies	Stocks			Ratio		(%)
Individual Stock	499	50,031	-0.083(-24.8)	0.799(133.0)	0.002(24.2)	26.23
ATMF IV						
S&P500 ATMF IV	499	50,031	-0.059(-17.5)	0.793(131.02)	0.002(12.7)	25.61
VIX	499	50,031	-0.054(-17.8)	0.794(131.1)	0.002(12.6)	25.60

Shows the linear regression: $C_{rx} = a_0 + a_1 \frac{r}{\sigma} + a_2 L + \varepsilon$ for S&P500 stocks (2005-2014), where C_{rx} is the correlation between returns at time, *t*, and volumes at time, *t*. *r* and σ are the returns and standard deviations of returns respectively, over the given period, *P* is a proxy measure for loss-aversion which can be the VIX, the ATMF IV of the individual stock or the ATMF IV of the S&P500. The t-values are shown in brackets. Results are separated into Panels A, B and C to reflect annual, quarterly and monthly observations respectively. The return-volume correlation is positively correlated with the Sharpe Ratio, the VIX and the ATMF implied volatilities of the S&P500 and the individual stock. This is seen at each observation frequency and is highly statistically significant. Results are consistent with the model predictions.

3.6 Conclusions

Using an analytical model, developed by Barberis and Xiong (2009), I showed that the contemporaneous return-volume is governed by the optimal trading strategy. Negative correlations were found to be associated with contrarian strategies and positive correlations with herding strategies. The optimal trading strategy was found to be governed by expected stock returns, the standard deviations of returns and investor preferences. These findings were verified empirically using a broad range of developed and emerging market stocks and the trading activity of institutional investors on S&P500 stocks.

The findings of this research support those of Avramov et al. (2006) in so far as identifying the importance of herding and contrarian strategies in the Leverage Effect. However, the manner in which these trading strategies manifest are different. They hypothesised that it was the interplay between the two strategies that generated the Leverage Effect, whereas this model suggests that the two strategies lead to different types of effect i.e. positive and negative Leverage Effect. In addition, I have provided a direct mechanism through which these strategies manifest and shown under what conditions they are optimal.

There are several features which do not make the Barberis and Xiong (2009) model directly applicable to the real world. Most importantly it assumes a complete market, which means that the entire price path is known and there is no friction to transacting. These assumptions are clearly unrealistic since in practice it is extremely difficult to predict the evolution of the stock price and in order to trade one must pay transaction fees. These assumptions are made to make the problem analytically tractable. The model also makes no predictions when the expected returns are below the risk-free return because in this scenario an investor would not invest in the stock. Barberis and Xiong (2009) also note several differences in the trading behaviour of this model compared to actual trading behaviour. Firstly, the trading behaviour in the model involves partial adjustments to risky asset holdings rather than selling entire positions. Secondly, assuming that the stock returns are binomially distributed leads investors to use leverage. The binomial distribution makes the model more tractable but it specifies a positive lower bound on the gross stock return which leads to aggressive allocations. They found that using a log normal return

distribution leads to far less leverage. Whilst these two factors may be a concern for retail investors they may actually be more appropriate for professional investors.

As I have shown that return-volume correlation is dependent upon an individual's preferences, it is highly likely that it is also affected by other behavioural biases. For instance, the 'House Money' Effect suggests that an individual's preferences may vary over time because investors are less concerned with losing profits than they are with losing initial capital. This means that after periods of positive returns, where they have amassed wealth, they will be more risk-seeking. Whereas, when they have endured periods of losses, future losses become more painful and so they become more risk-averse. Furthermore, investors may try to project current trends into the future (hindsight bias) and hence alter their trading behaviour accordingly; this would mean that return-volume correlation would lead returns. In Part II I move on to examine the Leverage Effect at the index level.

Part II

The Leverage Effect in Stock Indices

Chapter 4

Literature Review: The Index Level Leverage Effect

4.1 Introduction

This chapter examines the current literature on the Leverage Effect in stock indices. This research has historically taken two approaches. The first is a bottom up approach where stock level effects are aggregated to the index level and examined in the context of Black's 'Leverage Hypothesis' (Section 4.2) and Bouchaud et al. (2001) 'Retarded Volatility Hypothesis' (Section 4.4). The second is a top down approach, using the 'Volatility Feedback Hypothesis' (Section 4.3) and 'Collective Behaviour Hypothesis' (Section 4.6). Other effects that are also often associated with the index level Leverage Effect are trading volumes (Section 4.5) and stock correlation asymmetry (Section 4.7).

4.2 Leverage Hypothesis

Figlewski and Wang (2000) - studied the S&P 100 index (1977-1996), at the monthly and quarterly frequency, using the regression in Equation 1.7. From Table 4.1 it is evident that a_1 is far larger at the index level than at the stock level but the t-statistic (shown in brackets) shows that it is not statistically significant.

Table 4.1: Leverage Effect at the Stock and Index Levelsfrom Figlewski and Wang (2000)

Frequency	Index Level (a_1)	Stock Level (a_1)
Monthly	-0.452(-0.710)	-0.340(-14.218)
Quarterly	-0.644 (-1.206)	-0.254(-8409)

This compares the Leverage Effect at the index and stock levels as estimated by Figlewski and Wang (2000) using the regression in Equation 1.7. The t-values are shown in brackets. The results show that the Leverage Effect at the index level is far larger but less statistically significant.

An examination of the up and down market effects using the regression in Equation 1.8 shows that $a_1 > 0$ and $a_2 < 0$. This indicates that when the market falls, volatility increases but as the market rises volatility also increases, with the 'down market' effect dominating. Using the regressions in Equations 1.9 and 1.10 they also find that changes in volatility die out over time which is inconsistent with the premise that a permanent change in leverage should lead to a permanent change in volatility.

Daouk and Ng (2011) - studied equally weighted portfolios based on stocks in the S&P400, S&P500 and S&P600. They found that the Leverage Effect is significant for both the index and the portfolio; with the former being far larger. However, once the assets are unlevered, the effect in the index is only reduced by 11% but in the portfolios it is reduced by 77%. This implies that leverage is not responsible for the Leverage Effect observed at the index level. They also document a 'down market' effect for the index with volatility rising by 4.64% more when the market falls by 1% than when it rises by 1%.

Hasanhodzic and Lo (2011) - studied equally weighted portfolios. Using Equations 1.1 and 1.2 they find a strong inverse relationship between a firm's returns and the resulting change in volatility. They also find that the effect is larger for all equity financed firms than all debt financed firms and that the effect persists over time. However, they have not explicitly compared the aggregated data to that of the index level.

4.3 Volatility Feedback Hypothesis

As outlined in Chapter 1, this relies upon two properties. Firstly, volatility must be persistent so that a large realisation of news increases both current and future volatility and secondly, that there is a positive inter-temporal relation between expected return and conditional variance. Therefore, an increase in volatility leads to an increase in expected returns and hence a decrease in the current stock price. This dampens volatility in the case of good news and increases volatility in the case of bad news. However, research regarding the 'Volatility Feedback Hypothesis' at the index level has generated conflicting results with French et al. (1987) and Campbell and Hentschel (1992) finding that the relation between volatility and expected return is positive, whilst Turner et al. (1989), Glosten et al. (1993) and Nelson (1991) found that the relation is negative. Commenting on this conflict, Bekaert and Wu (2000) state that "often the coefficient linking volatility to returns is statistically insignificant" and that "if the relation between market conditional volatility and market expected return is not positive then the validity of the time varying risk premium is in doubt".

One of the main mechanisms for investigating the volatility feedback hypothesis has been through the use of pricing and equilibrium models. Campbell and Hentschel (1992) were the first to develop a formal model of volatility feedback. They modelled stock dividends as a quadratic GARCH process and linked the dividend volatility to returns; assuming a linear relationship. The model is able to generate asymmetric volatility, negative skewness and excess kurtosis.

Wu (2001) developed a feedback model that shows that both the Leverage Effect and Volatility Feedback are statistically significant but the Leverage Effect contributes more than twice as much to the negative correlation between return and volatility. However, whilst the Volatility Feedback only appears to play a minor role during stable market conditions it appears to have a big impact during volatile periods.

However, Bollerslev et al. (2006), who identified the Leverage Effect in high frequency data, notes that risk based explanations, such as those by Campbell and Hentschel (1992), are thought of as applying to much coarser time intervals; monthly or quarterly. Therefore it is not clear that a risk based explanation can adequately account for the Leverage Effect detected at the intraday level.

4.4 Retarded Volatility Hypothesis

Bouchaud et al. (2001) analysed seven international stock markets from the US, Europe and Asia (1990-2000), at the daily frequency, using the cross-correlation function (Equation 1.11). They found that there exists a different time decay for stock indices (approximately 10 days) than for individual stocks (approximately 50 days) and that the amplitude of the correlation is much stronger for indices. As at the stock level, they also state that the cross-correlation function is well fitted by an exponential curve. However, the Retarded Volatility Hypothesis (as outlined in Section 1.4) cannot hold, as the estimate $C_{rv}(\tau \rightarrow 0) = -2$, is out by more than an order of magnitude.

4.5 Volume Hypothesis

At the stock level I showed compelling evidence that the Leverage Effect is caused by trading volumes which is consistent with the research of Avramov et al. (2006). Gallant et al. (1992) proposed that trading volumes also explain the Leverage Effect at the index level. They studied the S&P composite index (1928-1985) at the daily frequency using a semi-nonparametric model (Gallant and Tauchen (1989) and Gallant and Tauchen (1993)) of the conditional joint density of returns and volumes. They found that:

1) The market is equally likely to rise or fall on heavy volumes

2) Days of high volume are associated with high volatility

3) The Leverage Effect is not apparent in the bivariate distribution (only the univariate distribution)

Hence they conclude that the Leverage Effect can be fully explained by trading volumes.

4.6 Models of Collective Behaviour

A number of authors have attempted to replicate stylised facts using models of collective behaviour. Donangelo et al. (2006) developed the Fear Factor Model which was able to replicate gain-loss asymmetry which refers to losses tending to have larger magnitudes than gains. In the Fear Factor Model investors synchronise their actions during falling markets i.e. there are stronger stock-stock correlations during falling markets than rising ones. This simulates risk-aversion, where the utility loss of negative returns is larger than the utility gain for positive returns. This model is supported by the empirical research of Balogh et al. (2010), who concluded that the observed correlation asymmetry suggests a constant fear factor among stockholders. Ahlgren et al. (2007) extend the Fear Factor model to develop the Frustration Governed Market Model which also replicates the Leverage Effect. This model allows for the stock dynamics to be controlled via a time-dependent stock volatility. The Frustration Governed Market Model is constructed in the following way:

The index value, I_t , may be constructed according to:

$$I_{t} = \frac{1}{N} \sum_{i=1}^{N} \exp[P_{i,t}]$$
(4.1)

where $P_{i,t}$ is the log price of stock *i* at time, *t* and *N* is the number of constituent stocks. All stocks are assumed to perform geometrical Brownian motions with common volatility, σ_t , and uncorrelated Gaussian distributed idiosyncrasies, $\varepsilon_{i,t}$.

$$P_{i,t} = P_{i,t-1} + \varepsilon_{i,t} \sigma_t \tag{4.2}$$

In a downward trending market, the investors may find it necessary to replace some of their investments in order to adapt to an unfavourable situation; depending on the magnitude of the move and their stop-loss strategy etc. One may say that during these periods of falling prices, investors get frustrated and as a consequence the volatility of the stocks, σ_t , increase. On the other hand, when the trend is upwards, there is little point in altering a profitable position. Hence 'excited frustrated states' are gradually relaxed by lowering,

 σ_t , towards a more fundamental long term level, σ_0 . The dynamics of the volatility, σ_t , are given by:

$$\frac{\partial \sigma_t}{\partial t} = -\frac{\sigma_t - \sigma_0}{\kappa} - A f_t r_t \tag{4.3}$$

where r_t is the index return at time, t, f_t is a dummy variable ($f_t = 1$ for $r_t < 0$ and $f_t = 0$ otherwise), κ is the characteristic volatility decay time and A is a positive amplitude.

4.7 Stock Correlation Hypothesis

Reigneron et al. (2011) state that "the volatility of an index in fact reflects both the volatility of the underlying single stocks and the average correlation between these stocks. The increased Leverage Effect for indices (as opposed to stocks) must therefore mean that both of the quantities are sensitive to a downward move of the market". Whilst Bekaert and Wu (2000) conclude that for Volatility Feedback Hypothesis to explain the Leverage Effect, negative shocks at the index level must lead to an increase in the conditional covariances.

4.7.1 Empirical Evidence

Ang and Chen (2002) conduct an empirical investigation of stock correlations and noted the following properties:

- 1. Asymmetries between upside and downside correlation exists between stocks across international stock-markets.
- 2. Correlation asymmetries are greater for extreme downward moves.
- 3. Smaller stocks exhibit greater correlation asymmetry than larger stocks.
- 4. Value stocks demonstrate more asymmetry than growth stocks.
- 5. Stocks that have experienced recent loses have higher asymmetry.

- 6. Higher risk stocks, as measured by their betas, have a smaller degree of asymmetry than traditional defensive stocks.
- 7. There is no correlation between leverage and correlation asymmetries (controlling the size).

Reigneron et al. (2011) decomposed the Leverage Effect into two contributions: one from the dependence of the average stock volatility on the past returns of the index and a second one describing the average correlation. They find that:

- 1. The two contributions to the index leverage are of the same order of magnitude.
- 2. The correlation effect is stronger at short time scales but decays faster than the volatility effect.
- 3. The sum of the two fitted exponentials reproduces satisfactorily the full Leverage Effect, although the latter is underestimated at short times.

These findings are consistent with those of Daouk and Ng (2011) who show that while stock level unlevered volatility stays the same in a down market, index level unlevered volatility is higher due to the higher covariance between stocks.

4.7.2 Relationship between the Index and Stock Levels

Having established that stock correlation asymmetry exists, it is important to understand why the effect emerges and how this stock level property relates to the index level. Kwon and Oh (2012) used (discrete) *TE* to examine the amount of information that was transmitted between stocks and the index. They found that more information is sent from the index to the stocks than vice versa and that developed markets exhibit a larger asymmetry than emerging markets. In another study, Pan and Sinha (2007) studied the National Stock Exchange (NSE) of India by means of Random Matrix Theory, and found strong stock correlations, as well as strong correlations between the stocks and a market mode (a 'Market Index' constructed according to the leading eigenvector of the correlation matrix). They propose that the existence of such strong correlations between the stocks and the market variable is a characteristic feature of emerging markets. However, in a study of the S&P500 and Tel Aviv markets, Shapira et al. (2009) show that the effect of the index is more intricate. They use partial correlations to identify the strongest mediating variable in an index. They find that the effect of the index is an order of magnitude larger than the effect of any other stock. They posit that the observed strong correlations between the various stocks are mainly a reflection of the correlation of each one with the index. They also showed that the index provided additional latent information in the correlation matrix and that this varies significantly over time (Kenett et al., 2010).

4.8 Summary

This review showed evidence of the Leverage Effect at the index level and showed that it is far stronger than at the stock level. It has been identified at the daily and monthly frequencies and even in high frequency data. Through the use of cross-correlation functions, Bouchaud et al. (2001), Ahlgren et al. (2007) and Bollerslev et al. (2006) also claim that returns affect volatility and not vice versa. Bouchaud et al. (2001) also showed that the effect decayed approximately exponentially over a period of around 10 days which is far less than the approximately 50 days they observed at the stock level.

It appears that the index level Leverage Effect is the combination of several factors. The first is the aggregated Leverage Effect present in the constituent stocks and the second is a feedback mechanism which originates at the index level and manifests as a change in the stock correlations. The stock correlation asymmetry is consistent with the 'Volatility Feedback Hypothesis' and 'Models of Collective Behaviour'. It also appears that trading volumes play an important role.

In the next chapter, I will examine the effect of stock correlation asymmetry and trading volumes on the Leverage Effect from an information theoretic perspective.

Chapter 5

An Information Theoretic Analysis of Index Returns, Volatility and Trading Volumes

5.1 Introduction

Chapter 4 examined the Leverage Effect at the index level. It showed that the 'Leverage Hypothesis' and the 'Retarded Volatility Model' were both unable to replicate the magnitude of the effect whilst the 'Volatility Feedback Hypothesis' only appeared significant during volatile periods. The most promising research is based upon trading volumes (Gallant et al., 1992), similar to the stock level, and stock correlation asymmetry (Reigneron et al., 2011). Stock correlation asymmetry posits that when the market declines a larger proportion of the stocks also decline than rise when the market rises. This leads to a magnification of the stock level Leverage Effect. In this chapter I examine the evidence for these hypotheses from an information theoretic perspective - the methodology is based upon Chapter 2 - and find that both play an important role in the Leverage Effect.

The chapter initially details the data and model calibration in Section 5.2 further information on the methods can be found in Section 2.2. The results are shown in Section 5.3 with the conclusions given in Section 5.4.

5.2 Methodology

5.2.1 The Data

The data is sourced from Bloomberg and covers daily returns and volumes for developed stock indices covering the period 1980-2012. The indices include Australia (ASX), Canada (TSX), France (CAC), Germany (DAX), Holland (AEX), Hong Kong (HSI), Italy (MIB), Japan (NKY), South Korea (KOSPI), Spain (IBEX), Sweden (OMX), Switzerland (SMI), UK (FTSE), USA (DJIA, NDX and SPX). The trading volume metric is given as the sum of the daily volumes of the constituent stocks. Hence, the following research does not consider the impact of derivatives such as futures and forwards on the underlying index properties. Derivatives volumes are likely to have an impact on the underlying index returns and volatility as groups of stocks are often traded against futures and other derivatives. This is certainly an interesting area for future research.

All statistics and results have been calculated at the individual index level and then averaged to give a mean value for all indices. The significance levels have been estimated by calculating the various measures using surrogate data sets with similar statistical properties but without the inter-relationships; this is consistent with similar research.

5.2.2 **Calibrating the Model**

First, let me define the variables with which I will conduct the analysis (these are at the daily frequency):

Returns:

$$r_t = \ln[I_t] - \ln[I_{t-1}] \tag{5.1}$$

where I_t is the index value at time, t. Normalised Returns:

> $\hat{r}_t = \frac{r_t}{\sigma_r}$ (5.2)
where σ_r is the standard deviation of the returns.

Volatility:

$$v_t = r_t^2 \tag{5.3}$$

Normalised Volatility:

$$\hat{v_t} = \frac{v_t}{\sigma_v} \tag{5.4}$$

where σ_v is the standard deviation of the volatility. Normalised Volume:

$$\hat{x}_t = \frac{\ln\left(x_t\right) - \ln\left(\bar{x}\right)}{\sigma_x} \tag{5.5}$$

where x_t is the index volume at time, t, and σ_x is the standard deviation of the volumes.

Here I have normalised variables by dividing by the standard deviation of the whole sample of the stock data. The purpose of this is to set each stock to unit variance so that they can easily be compared and to improve the convergence of the estimators. By taking this approach it may bias some techniques because one is influencing historic returns by future volatility. However, the *MI* estimator should in principle be independent of scale factors.

The daily correlation between the constituent stocks, C_t , is calculated as:

For positive index returns:

$$C_t = \frac{1}{N} \sum_{i=1}^N \Lambda_i \tag{5.6}$$

For negative index returns:

$$C_t = 1 - \frac{1}{N} \sum_{i=1}^N \Lambda_i \tag{5.7}$$

where

$$\Lambda_i = \begin{cases} 1 & r_i \ge 0 \\ & for \\ 0 & r_i < 0 \end{cases}$$
(5.8)

N is the total number of stocks in the index and r_i is the daily stock return. Therefore, one is simply taking the proportion of stocks that move in the same direction as the index on any given day.

The results are given using the stretched exponential function given in Equation 2.18.

In order to calibrate the algorithm it is necessary to ascertain the number of nearest neighbours, k, which is used to estimate the probability density and the number of observations required to generate stable results. Figure 5.1 shows the $MI(\hat{r}_t, \hat{v}_{t+1}), MI(\hat{r}_t, \hat{x}_{t+1})$ and $MI(\hat{x}_t, \hat{v}_{t+1})$. Figure 5.1 (Top) shows that the *MI* converges as the number of nearest neighbours, k, increases for each of the measures. It indicates that the measures have largely converged for $k \approx 75$. Figure 5.1 (Bottom) shows the *MI* for varying data lengths from 250 - 2500 days; k = 75. The *MI* appears stable for data lengths greater than 1500 - 2500 which implies that over these periods the mean *MI* is stationary; this is a requirement to estimate stable and consistent results. Consequently, in the following analysis I will use k = 75 and exclude indices with less than 1500 observations.



Figure 5.1: Mutual Information Calibration for Indices at the Daily Frequency

(Top) Shows the mean *MI* for developed market indices as a function of the number of nearest neighbours, *k*. The blue stars with dashed line is the $MI(\hat{r}_t, \hat{v}_{t+1})$, the red diamonds with solid line is the $MI(\hat{r}_t, \hat{x}_{t+1})$ and the green squares with dot-dashed line is $MI(\hat{x}_t, \hat{v}_{t+1})$. These are given with associated one standard errors and exponential curve fits. The results appear to show that the *MI* converges as the number of nearest neighbours increases and that $k \approx 75$ should be sufficient to produce stable results. (Bottom) Shows the mean *MI* for developed market indices across a range of data lengths (days). The blue stars with dashed line is the $MI(\hat{r}_t, \hat{v}_{t+1})$, the red diamonds with solid line is the $MI(\hat{r}_t, \hat{x}_{t+1})$ and the green squares with dot-dashed line is the $MI(\hat{r}_t, \hat{v}_{t+1})$. These are given with associated one standard errors and k = 75. The results indicate that the *MI* does not vary considerably for observation periods between 1500 - 2500 days. This suggests that the mean *MI* is stationary over this time period.

5.3 Results

5.3.1 Persistence of Returns, Volatility and Volumes

The first step to understanding the properties of index returns, volatility and trading volumes is to analyse their auto-mutual information and auto-covariance functions. These functions indicate how the variables persist over time. Figure 5.2 shows the auto-mutual information (Top) and auto-covariance functions (Bottom) for returns, volatility and volumes for developed market indices. It shows that the *MI* for returns and volatility is statistically significant for $\tau < 20$ days but the *MI* for volume is significantly larger and persists for $\tau > 30$ days. The notable difference from the stock level is that the auto-mutual information for volatility is smaller and consequently becomes statistically insignificant in a shorter period of time. The auto-covariance functions for volumes and volatility also show that they are highly persistent and statistically significant; with the former being of a larger magnitude. The notable difference is the lack of auto-covariance in returns; this has been documented previously (Cont, 2001). This difference could be an indication that the lower order moments are arbitraged out in the market. This is important as it could lead to spurious results when using linear causality analysis.

- Returns, Volatility and Volumes all display Auto-Information



Figure 5.2: Auto-Mutual Information and Auto-Covariance Functions for Indices at the Daily Frequency

Shows the auto-mutual information (Top) and normalised auto-covariance (Bottom) functions. They are calculated for returns (blue stars with dashed lines), volatility (green squares with dot-dashed lines) and volumes (red diamonds with solid lines). These are given with associated one standard errors, exponential curve fits and 95% Significance Levels (black dashed lines). The *MI* is statistically significant for returns and volatility for approximately 20 days whilst the *MI* for volumes significantly larger and is persistent for longer than 30 days. The auto-covariance function exhibits many of the same properties as the auto-mutual information function except that the auto-covariance for returns is not statistically significant at any time horizon.

5.3.2 Returns and Volumes

Figure 5.3 shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for returns and volume. It shows that both the $MI(\hat{r}_t, \hat{x}_{t+\tau})$ and $MI(\hat{x}_t, \hat{r}_{t+\tau})$ are statistically significant and persistent. $MI(\hat{r}_t, \hat{x}_{t+1}) > MI(\hat{x}_t, \hat{r}_{t+1})$ but equal for $\tau > 1$. Hence the *MI* implies a bi-directional information flow between returns and volumes. These results contrast with those of the cross-covariance function which shows that the covariance is not statistically significant in either direction. This also differs from the stock level which shows statistically significant covariances for $Cov(r_t, x_{t+\tau})$.

Figure 5.4 (Top) shows the partial cross-mutual information function for returns and volume where I have controlled for auto-information. It shows that only the $MI(\hat{x}_t, \hat{r}_{t+\tau}|Z)$ for $1 \leq \tau \leq 3$ is statistically significant at the 95% confidence level. This implies that volumes cause returns and not vice versa. In addition, since the *PMI* decays far more quickly than the *MI* it also implies that the persistence is due to auto-information. The results for the *TE* (Figure 5.4 (Bottom)) shows that $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau}) > TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ for $\tau \leq 100$. However, they only support the results of the *PMI* for $\tau < 4$ because beyond this point, $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ is statistically significant. This indicates that the *PMI* may be unduly influenced by the persistence of the linear correlations. The *TE* implies bi-directional (Granger) causality between volumes and returns with volumes dominating.

- Persistence in the return-volume relation is driven by auto-information

- There is bi-directional (Granger) causality between returns and volumes with volumes dominating



Figure 5.3: Cross-Mutual Information and Cross-Covariance Functions for Index Returns and Volume at the Daily Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions between returns and volume. The blue squares with solid lines show the $MI(\hat{x}_t, \hat{r}_{t+\tau})$ and $Corr(x_t, r_{t+\tau})$. The green diamonds with dot-dashed lines show the $MI(\hat{r}_t, \hat{x}_{t+\tau})$ and $Corr(r_t, x_{t+\tau})$. These are given with associated one standard errors, stretched exponential curve fits and 95% Significance Levels (dashed black lines). The MI is statistically significant and persistent in both directions with $MI(\hat{r}_t, \hat{x}_{t+1}) > MI(\hat{x}_t, \hat{r}_{t+1})$ but approximately equal for $\tau > 1$. This indicates bi-directional causality between returns and volumes. However, the cross-covariance function shows no statistically significant covariance in either direction.



Figure 5.4: Partial Cross-Mutual Information and Cross-Transfer Entropy Functions for Index Returns and Volume at the Daily Frequency

(Top) Shows the partial cross-mutual information function for returns and volume, where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{x}_{t+\tau}|Z)$ where Z = $[\hat{j}_t,...,\hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised returns. It is given with associated one standard errors, stretched exponential curve fits and a 95% Significance Level (dashed black line). It shows that only the $MI(\hat{x}_t, \hat{r}_{t+\tau}|Z)$ is statistically significant at the 95% confidence level for $\tau < 4$. Hence, the PMI implies that volumes cause returns and not vice versa. In addition, since both PMI decay more quickly than the MI it also implies that the persistence is due to auto-information. (Bottom) Shows the cross-transfer entropy function for returns and volumes. The green diamonds represent $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ and the blue squares represent the $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ and the dashed blue line for $TE(\hat{x}_t \to \hat{r}_{t+\tau})$. The *TE* results show that $TE(\hat{x}_t \to \hat{r}_{t+\tau}) > TE(\hat{r}_t \to \hat{x}_{t+\tau})$ for all $\tau \lesssim 100$. However, they only support the *PMI* results for $\tau < 4$ because beyond this point, $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ is statistically significant. This indicates that the PMI may be unduly influenced by the persistence in the linear correlations. The TE implies bi-directional (Granger) causality between volumes and returns with volumes dominating.

5.3.3 Volume and Volatility

Figure 5.5 shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for volumes and volatility. Again Figures 5.3 and 5.5 are similar because the volatility is simply the square of the returns, so the *MI* struggles to differentiate between the two variables. $MI(\hat{v}_t, \hat{x}_{t+\tau})$ and $MI(\hat{x}_t, \hat{v}_{t+\tau})$ are both statistically significant and persistent with $MI(\hat{v}_t, \hat{x}_{t+1}) > MI(\hat{x}_t, \hat{v}_{t+1})$ but for $\tau > 1$ they are of equal magnitudes. This implies a bi-directional information flow between volumes and volatility. The results of the cross-mutual information are consistent with those of the cross-covariance which show that the $Cov(v_t, x_{t+\tau})$ and $Cov(x_t, v_{t+\tau})$ are both positive, of equal magnitude and statistically significant for $\tau > 30$.

Figure 5.6 (Top) shows the partial cross-mutual information function for volumes and volatility, where I have controlled for auto-information. It shows that the *MI* is only statistically significant for $MI(\hat{x}_t, \hat{v}_{t+\tau} \mid Z)$ for $\tau < 4$. Hence the *PMI* implies that volumes cause volatility and not vice versa. Again, since the *PMI* decays more quickly than the *MI* it also implies that the persistence is due to auto-information. The cross-transfer entropy function (Figure 5.6 (Bottom)), indicates that the *PMI* results are again likely the result of persistent linear correlations because the $TE(\hat{v}_t \rightarrow \hat{x}_{t+\tau})$ is statistically significant for $\tau \gtrsim 20$. Therefore the *TE* results indicate bi-directional (Granger) causality between volumes and volatility with volumes dominating.

- Persistence in the volume-volatility relation is driven by auto-information

- There is bi-directional (Granger) causality between volatility and volumes with volumes dominating



Figure 5.5: Cross-Mutual Information and Cross-Covariance Functions for Index Volatility and Volume at the Daily Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for volumes and volatility. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau})$ and the $Cov(v_t, x_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau})$ and the $Cov(x_t, v_{t+\tau})$. Both graphs are given with associated one standard errors, stretched exponential curve fits and 95% Significance Levels (black dashed lines). The MI is statistically significant and persistent in both directions with $MI(\hat{v}_t, \hat{x}_{t+1}) > MI(\hat{x}_t, \hat{v}_{t+1})$ but approximately equal for $\tau > 1$. This indicates bi-directional causality between volume and volatility. This is consistent with the results of the cross-covariance function which shows that $Cov(v_t, x_{t+\tau})$ and $Cov(x_t, v_{t+\tau})$ are positive, of equal magnitude and statistically significant for $\tau > 30$.



Figure 5.6: Partial Cross-Mutual Information and Cross-Transfer Entropy Functions for Index Volatility and Volume at the Daily Frequency

(Top) Shows the partial cross-mutual information function for volumes and volatility where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volatility. It is given with associated one standard errors, stretched exponential curve fits and a 95% Significance Level (black dashed line). The MI is only statistically significant for $MI(\hat{x}_t, \hat{v}_{t+\tau} | Z)$ for $\tau < 4$ which indicates that volumes cause volatility and not vice versa. Since the PMI decays more quickly than the MI it also implies that the persistence is due to auto-information. (Bottom) Shows the cross-transfer entropy function for volumes and volatility. The green diamonds represent $TE(\hat{v}_t \rightarrow \hat{x}_{t+\tau})$ and the blue squares represent $TE(\hat{x}_t \rightarrow \hat{v}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels where the green dot-dashed line is for the $TE(\hat{v}_t \rightarrow \hat{x}_{t+\tau})$ and the dashed blue line is for the $TE(\hat{x}_t \rightarrow \hat{v}_{t+\tau})$. This shows that $TE(\hat{x}_t \rightarrow \hat{v}_{t+\tau}) > TE(\hat{v}_t \rightarrow \hat{x}_{t+\tau})$ for all τ . However, they only support the results of the PMI for $\tau \leq 20$ because beyond this point, $TE(\hat{v}_t \rightarrow \hat{x}_{t+\tau})$ is statistically significant. This indicates that the PMI may be unduly influenced by the persistence of the linear correlations. The TE implies bi-directional (Granger) causality between volumes and volatility with volumes dominating.

5.3.4 Returns and Volatility

Figure 5.7 shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for returns and volatility. The cross-mutual information function shows that the $MI(\hat{r}_t, \hat{v}_{t+\tau}) > MI(\hat{v}_t, \hat{r}_{t+\tau})$ for $\tau < 3$ and statistically significant. For $\tau \ge 3 MI(\hat{r}_t, \hat{v}_{t+\tau}) \approx MI(\hat{v}_t, \hat{r}_{t+\tau})$ and neither are statistically significant beyond $\tau \ge 25$. The *MI* is of similar magnitude to the stock level but the decay rate is faster. Unusually, the *MI* for $\tau = 2$ is greater than for $\tau = 1$. The *MI* function implies a bi-directional information flow between returns and volatility. The Leverage Effect is clearly identifiable in the cross-covariance function with structure for $Cov(r_t, v_{t+\tau})$; this is statistically significant for $\tau \le 15$. Some authors, such as Bouchaud et al. (2001), have found this sufficient to imply causation from returns to volatility. As documented previously the magnitude of the Leverage Effect at the index level is twice that at the stock level.

Figure 5.8 (Top) shows the partial cross-mutual information function for returns and volatility where I have controlled for auto-information. The *PMI* estimate is overstated because it is not possible to control for the auto-information from the variables when $\tau = 1$ because the volatility is directly calculated from the returns and hence controlling for one indirectly controls for the other. The *PMI* shows bi-directional information and that the persistence is caused by auto-information due to the faster decay. The *TE* also shows bi-directional information flow (Figure 5.8 (Bottom)).

- Persistence in the return-volatility relation is driven by auto-information

- There is bi-directional (Granger) causality between returns and volatility



Figure 5.7: Cross-Mutual Information and Cross-Covariance Functions for Index Returns and Volatility at the Daily Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for returns and volatility. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau})$ and the $Cov(r_t, v_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau})$ and the $Cov(v_t, r_{t+\tau})$. Both graphs are given with associated one standard errors, stretched exponential curve fits and 95% Significance Levels (black dashed lines). The cross-mutual information function shows that $MI(\hat{r}_t, \hat{v}_{t+\tau})$ and $MI(\hat{v}_t, \hat{r}_{t+\tau})$ are both statistically significant. It also shows that $MI(\hat{r}_t, \hat{v}_{t+\tau}) > MI(\hat{v}_t, \hat{r}_{t+\tau})$, for $\tau < 3$. It implies a bi-directional information flow between returns and volatility. The Leverage Effect is clearly identifiable in the crosscovariance function which only show structure for $Cov(r_t, v_{t+\tau})$; which is statistically significant for $\tau \leq 15$. Some authors have found this sufficient to imply causation from returns to volatility. The magnitude of the Leverage Effect is twice that at the stock level whilst the MI is of similar magnitude.



Figure 5.8: Partial Cross-Mutual Information and Cross-Transfer Entropy Functions for Index Returns and Volatility at the Daily Frequency

(Top) Shows the partial cross-mutual information function for returns and volatility where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 2$ and \hat{j} are the normalised volatilities. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 2$ and \hat{j} are the normalised returns. It is given with associated one standard errors, stretched exponential curve fits and a 95% Significance Level (black dashed lines). The *PMI* indicates bi-directional information flow and that the persistence is due to autoinformation. (Bottom) Shows the cross-transfer entropy function for returns and volatility. The green diamonds represent the $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the blue squares represent $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the dashed blue line for $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. The *TE* also shows bi-directional information flow. This implies bi-directional (Granger) causality between volatility and returns with volatility dominating. I now examine the impact of trading volumes and correlation asymmetry on the Leverage Effect using *PMI* (Table 5.1). Here I consider the 1 and 2 day time lags because the *MI* is actually larger at the 2 day lag. A number of interesting features are identifiable. Firstly, the *MI* is of a similar magnitude to the stock level (Table 2.1). This may seem surprising since it has been widely documented (using correlations/covariances) that the Leverage Effect is far larger at the index level. Secondly, controlling for trading volumes has little or no effect on the *MI* between returns and volatility. Thirdly, controlling for correlation asymmetry actually increases the *MI* by two orders of magnitude. Fourthly, when controlling for trading volumes and correlation asymmetry the *MI* increases by only an order of magnitude.

These results imply that the *MI* function simply sees the index level Leverage Effect as an aggregation of the stock level Leverage Effect. However, controlling for correlation asymmetry removes the diversification effect revealing a much larger Leverage Effect. Then trading volumes also become important, accounting for 61.48% of the *MI* between returns and volatility. Unfortunately, the *MI* remains statistically significant but we cannot control for individual stock effects due to the high dimensionality of the problem.

- Correlation asymmetry acts as a dampening mechanism for the MI between returns and volatility

- Trading volumes are a driver for the return-volatility relation

Asymmetry $MI(\hat{r},\hat{v})$ $MI(\hat{r}, \hat{v} \mid \hat{x})$ $MI(\hat{r}, \hat{v} \mid \hat{c},)$ $MI(\hat{r},\hat{v}\mid\hat{x},\hat{c},)$ 0.006 ± 0.001 0.007 ± 0.001 0.135 ± 0.020 0.052 ± 0.009 (r_t, v_{t+1}) (0.005)(0.004)(<0.001) (<0.001) 0.003 ± 0.001 0.005 ± 0.001 0.134 ± 0.020 0.051 ± 0.009 (v_t, r_{t+1}) (0.112)(0.016)(<0.001) (<0.001)

 0.128 ± 0.020

(<0.001)

 0.126 ± 0.020

(<0.001)

 0.054 ± 0.010

(<0.001)

 0.053 ± 0.010

(<0.001)

 0.012 ± 0.001

(<0.001)

 0.010 ± 0.001

(0.001)

 (r_t, v_{t+2})

 (v_t, r_{t+2})

 0.012 ± 0.001

(<0.001)

 0.008 ± 0.001

(0.003)

Table 5.1: Mutual Information for the Index Level Leverage Effect Controlling for Trading Volumes and Correlation

Shows the *MI* between returns and volatility at the 1 and 2 day time lags for developed market indices. The p-values are shown in brackets. Controlling for trading volumes alone has no effect on the MI between returns and volatility but correlation asymmetry is shown to dampen the MI by two orders of magnitude. Controlling for correlation asymmetry and trading volumes increases the MI between returns and volatility by an order of magnitude. These results indicate that correlation asymmetry actually dampens the MI between returns and volatility but once this has been accounted for trading volumes become an important driver of the Leverage Effect.

5.4 Conclusions

In this chapter I identified the index level Leverage Effect in the cross-covariance function and showed the magnitude to be twice as large as at the stock level; consistent with previous research. However, the $MI(\hat{r}, \hat{v})$ was found to be similar to that at the stock level which indicated that in general it could be considered as the average of the stock level effect. The reason for this discrepancy appears to be because correlation asymmetry acts as a diversification factor and is a suppressing variable. Having controlled for this diversification effect, trading volumes then account for 62% of the *MI* between returns and volatility.

The chapter also produced a number of stylised facts, from an information theoretic perspective, which may give insights into the functioning of the financial markets:

1) Returns, volatility and volumes all display auto-information.

2) There is bi-directional (Granger) causality between volumes and stock returns but volumes dominate.

3) There is bi-directional (Granger) causality between returns and volatility.

4) There is bi-directional (Granger) causality between volumes and volatility but volumes dominate.

5) The persistence in the relationships between returns, volatility and volumes are driven by auto-information.

6) Appendix A shows that at the weekly frequency volumes (Granger) cause returns and volatility whilst volatility (Granger) cause returns.

These results are also found to be generally consistent with emerging market indices (Appendix C).

Chapters 5 and 3 have highlighted the importance of time dependence in the returnvolume correlation at the stock level and the Leverage Effect at the index level respectively. In the next chapter I examine the time dependence of these effects and the relationship between the stock and index levels.

Part III

Time Variation of the Leverage Effect

Chapter 6

Time-Series Analysis of the Leverage Effect

6.1 Introduction

In Parts I and II, I examined the Leverage Effect at the index and stock levels and associated effects such as stock correlation asymmetry and return-volume correlation. This examination suggests that these effects may in fact be time dependent. For example, in Chapter 3 I showed that return-volume correlation is governed by the optimal trading strategy. This in turn is governed by expected returns, the standard deviation of returns and investors risk preferences, all of which vary with time. Whilst in Chapter 5 I showed that the *MI* in the index level Leverage Effect is suppressed by stock correlation asymmetry. This is consistent with previous findings that 'The Volatility Feedback Hypothesis' is important in periods of distress but less so during quiet periods. In this chapter I examine how these effects vary over time and how they are related. Surprisingly there has been little previous research in this area. I find that all the effects are time dependent and coherent. It is shown that return-volume correlation leads the stock and index level Leverage Effects which is consistent with the notion that the Leverage Effect is driven by trading activity. It also shows that the stock and index level Leverage Effects and stock correlation asymmetry all move in phase during volatile periods when there is also a larger information transfer from the index to the stock level. This indicates that during these periods investors are more focussed on systemic factors.

The chapter will initially detail the data and methods used in Section 6.2; see Chapter 2.2 for details on TE. It then proceeds to present an empirical investigation of the time variation of the Leverage Effect at the stock and index levels in Section 6.3. Finally, I summarise my findings in Section 6.4.

6.2 Methodology

6.2.1 The Data

The data is sourced from Bloomberg and covers daily stock returns and volumes for 488 stocks from the S&P500 during the period 2000-2012. The stocks have a daily mean return of 0.0410%, a daily mean standard deviation of 0.0259 and a mean daily volume of 2.2493M shares. Since the data represents the constituent stocks as of 2012, the data set is prone to survivorship bias. Stocks that have ceased to trade or have dropped out of the index over this time may display different dynamics/relationships hence this research makes no statement about these stocks. This may be an interesting area of future research.

The data has been segmented into annual periods (252 day) and calculated on a weekly (5 day) rolling basis. It is necessary to use over-lapping time windows in order to produce enough data points to analyse but this does induce correlations/coherence between the time points. The stock level results are calculated at the individual stock level and then averaged to give a value for the overall index and associated standard errors. The index levels are shown as calculated without error estimates since there is only one realisation.

6.2.2 What are Wavelets?

Wavelets are mathematical functions used to represent data or other functions in a similar way to sine and cosine functions in Fourier Analysis. However, unlike sine and cosine functions, wavelets are 'local' which means they can examine signals at different resolutions or scales. This is particularly useful in many real world systems which are often non-stationary and require analysis at the local scale. To conduct wavelet analysis it is necessary to choose a prototype wavelet, often called a 'mother' wavelet. This 'mother' wavelet is then rescaled to examine different properties of the signal. A high frequency version is used for temporal analysis and a low frequency version for frequency analysis. The choice of 'mother' wavelet is a challenging one and often resorts to a trial and error approach in determining which wavelet shape most accurately reflects your data; this is generally done by visually examining the wavelet transform (discussed below). In the following analysis I use the cgau3 wavelet (Mathworks, 2014).

Dilations and translations of the 'mother' function, $\Phi(x)$, define an orthogonal basis, our wavelet basis:

$$\Phi_{(a,b)}(x) = 2^{-\frac{a}{2}} \Phi\left(2^{-a} x - b\right)$$
(6.1)

where x is a real valued time-series, a is the scale index and indicates the wavelet's width and b is the location index which gives the position. To span the data domain at different resolutions, the analysing wavelet is used in a scaling equation:

$$W(x) = \sum_{k=-1}^{N-2} (-1)^k c_{k+1} \Phi(2x+k)$$
(6.2)

where W(x) is the scaling function for the 'mother' function, Φ , and c_k are the wavelet coefficients.

Wavelet coefficients must satisfy linear and quadratic constraints of the form:

$$\sum_{k=0}^{N-1} c_k = 2 \tag{6.3}$$

$$\sum_{k=0}^{N-1} c_{kc_{k+2b}} = 2\delta_{b,0} \tag{6.4}$$

where δ is the delta function and *b* is the location index. The original signal or function can then be represented as a wavelet expansion using coefficients in a linear combination of the wavelet functions. The wavelet coefficients act as a filter and act in two ways, one bringing out the detail in the signal and the second smoothing the signal. Apply the coefficients to a signal using a transformation matrix results in a wavelet transform denoted by $C_x(a,b)$.

We can use the wavelet transform to examine the relationship between multiple timeseries using wavelet coherence, where one can calculate the modulus and phase angles of the coherence of the two signals. This is calculated via the wavelet cross-spectrum which is given by:

$$C_{xy}(a,b) = \overline{C_x(a,b)}C_y(a,b)$$
(6.5)

where \bar{z} denotes the complex conjugate of z.

This empirical wavelet coherence is then given as:

$$Coherence = \frac{K(C_{xy}(a,b))}{\sqrt{K(|C_x(a,b)|)^2}\sqrt{K(|C_y(a,b)|)^2}}$$
(6.6)

where *K* is a smoothing kernel.

For a further discussion on wavelets I refer the reader to Graps (1995) from which these notes were taken and Torrence and Campo (1998).

6.3 Results

6.3.1 Stock Level Effects

In Chapter 3 I showed that return-volume correlation is governed by the expected stock returns and the standard deviation of returns. Figure 6.1 shows how the Sharpe Ratio (μ/σ) and the return-volume correlation vary over time. The middle and bottom plots represent the modulus and phase angle of the wavelet coherence respectively. The scale parameter, *a*, is shown on the vertical axes (in weeks) and the location parameter, *b*, is on the horizontal axes. To interpret the modulus we can see that it is red across both scale and location which means the signals are strongly coherent; this is probably due to the overlapping time windows. The phase angle is far more interesting showing a range of

negative (blue) and positive (yellow/red) regions across both scale and location. It appears to show that return-volume correlation lead the Sharpe Ratio prior to the 2008 financial crisis this is observed by the blue regions for at 20-40 week wavelengths. This is consistent with using expected returns, but lags after the crisis. However, after the crisis the relationship reversed; identified by the yellow regions at 15-30 week wavelengths. This reversal could indicate that investors are exposed to hindsight bias where they are extrapolating past returns into the future; in this case they have been influenced by the financial crisis. The return-volume correlation was also found to be dependent upon an individual's preferences which are also time dependent. The 'House Money' Effect suggests that an individual's preferences may vary over time because investors are less concerned with losing profits than they are with losing initial capital. This means that after periods of positive returns, where they have amassed wealth, they will be more risk-seeking. Whereas, when they have endured periods of losses, future losses become more painful and so they become more risk-averse. However, the 'House Money' Effect it is not readily apparent in Figure 6.1. This may be because it is not a factor or because the effect of the Sharpe Ratio dominates or because the required frequency is below the annual level.



Figure 6.1: Time Variation of Return-Volume Correlation and the Sharpe Ratio

Time Variation of the Sharpe Ratio and Return-Volume Correlation. (Top) shows that the Sharpe Ratio (blue squares with dashed line) and the Leverage Effect (green diamonds with solid line) are generally both negative. (Middle) shows that the signals are coherent (red regions). (Bottom) shows that return-volume correlation tends to lead the Sharpe Ratio (negative regions) prior to the 2008 financial crisis but subsequently lags behind (positive regions).

Figure 6.2 examines the time varying relationship between the stock level Leverage Effect and return-volume correlation. It shows that they are both generally negative and coherent with return-volume correlation leading the Leverage Effect. This is consistent with the notion that trading volumes drive the stock level Leverage Effect. During the financial crisis they move out of phase with the Leverage Effect becoming more positive and the return-volume correlation becoming more negative.



Figure 6.2: Time Variation of the Stock Level Leverage Effect and Return-Volume Correlation

Time Variation of the Stock Level Leverage Effect and Return-Volume Correlation. (Top) shows that the return-volume correlation (blue squares with solid line) and the Leverage Effect (green diamonds with dashed line) are generally both negative. (Middle) shows that the signals are coherent (red regions). (Bottom) shows that return-volume correlation tends to lead the Leverage Effect (positive regions) especially prior to and during the 2008 financial crisis.

6.3.2 Index Level Effects

Figure 6.3 examines the time varying relationship between the index level Leverage Effect and return-volume correlation. It shows that they are both generally negative and coherent with return-volume correlation leading the Leverage Effect. This is consistent with the notion that trading volumes drive the index level Leverage Effect. It is again identifiable that they move out of phase during the financial crisis where the Leverage Effect becomes more positive and the return-volume correlation becomes more negative.



Figure 6.3: Time Variation of the Index Level Leverage Effect and Return-Volume Correlation

Time Variation of the Index Level Leverage Effect and Return-Volume Correlation. (Top) shows that the return-volume correlation (blue squares with solid line) and the Leverage Effect (green diamonds with dashed line) are generally both negative. However, they depart markedly around the time of the 2008 financial crisis when the return-volume correlation becomes more negative whilst the Leverage Effect actually becomes positive. (Middle) shows that the signals are coherent (red regions). (Bottom) shows that return-volume correlation tends to lead the Leverage Effect (positive regions) especially prior to and during the 2008 financial crisis.

6.3.3 Index and Stock Level Effects

Figure 6.4 examines the stock and index level Leverage Effects. It is apparent that they are coherent but the index level effect is far larger. It can also be seen that the stock level effect leads the index level until the 2008 financial crisis when they move in phase. Figures 6.5 and 6.6 show that the Leverage Effects lead correlation asymmetry until the 2008 financial crisis when they all move in phase. This regime change can clearly be seen in Figure 6.7 where there is a significant increase in the TE from index returns to stock returns and volatility after the crisis. This could indicate that the investors were more focussed on macro/systemic risks.



Figure 6.4: Time Variation of the Index and Stock Level Leverage Effects

Time Variation of the Index and Stock Level Leverage Effects. (Top) shows the index level Leverage Effect (blue squares with solid line) and the stock Leverage Effect (green diamonds with dashed line) are generally both negative but the index level effect is far larger. (Middle) shows that the signals are coherent (red regions). (Bottom) shows that the stock level Leverage Effect (positive regions) tends to to lead the index level effect except around the 2008 financial crisis when they move in phase.





Time Variation of the Index Level Leverage Effect and Correlation Asymmetry. (Top) shows that the correlation asymmetry (blue squares with solid line) and the Leverage Effect (green diamonds with dashed line) are generally both negative except around the time of the 2008 financial crisis when they both turn positive. (Middle) shows that the signals are coherent (red regions). (Bottom) the results show that generally the Leverage Effect leads the correlation asymmetry (positive regions) but around the time of the 2008 financial crisis they move in phase.





Time Variation of the Stock Level Leverage Effect and Correlation Asymmetry. (Top) shows that the correlation asymmetry (blue squares with solid line) and the Leverage Effect (green diamonds with dashed line) are generally both negative except around the time of the 2008 financial crisis when the correlation asymmetry turns positive. It is also clear that the correlation asymmetry is far larger than the stock level Leverage Effect. (Middle) shows that the signals are coherent (red regions). (Bottom) the results show that the stock level Leverage Effect generally leads the correlation asymmetry (positive regions) except around the time of the 2008 financial crisis when they move in phase.



Figure 6.7: Time Variation of the Transfer Entropy from the Index to Stock Level

Transfer Entropy between Index and Stock Levels. Index returns appear to transmit a similar amount of information to stock returns, $TE(I_r \rightarrow S_r)$ (red diamonds with dot-dashed line), as to stock volatility $TE(I_r \rightarrow S_v)$ (blue squares with a solid line). The index returns do not appear to transmit significant information to stock volumes as evidenced by $TE(I_r \rightarrow S_x)$ (green pentagrams with dashed line). The information transfer appears far higher subsequent to the 2008 financial crisis.

6.4 Conclusions

In this chapter I have shown that the Leverage Effects, return-volume correlation and correlation asymmetry all vary with time. Return-volume correlation was found to lead the Leverage Effect at the index and stock levels which is consistent with the notion that the Leverage Effect is driven by trading activity. It also showed that the stock level Leverage Effect leads the index Leverage Effect and both lead the correlation asymmetry except during volatile periods, such as the financial crisis, when they all move in phase.

The investigation also uncovered marked differences in the market dynamics prior to and after the 2008 financial crisis. It was found that subsequent to the 2008 financial crisis, stock returns and volatility were much more heavily influenced by index returns. This indicates that investors were more focussed on systemic risks. This is probably why the stock and index level Leverage Effects move in phase during these periods. It was also found that return-volume correlation led the Sharpe Ratio prior to the crisis which is consistent with the model predictions in Chapter 3 that suggested investors use expected stock returns and standard deviations. However, after the crisis this relationship flipped which could indicate that investors are exposed to hindsight bias where they are extrapolating past returns into the future. In this case they may have been 'fearful' after the crisis.

Part IV

The Leverage Effect in Stock Options

Chapter 7

Literature Review: The Implied Leverage Effect

Options are a type of financial instrument known as a derivative. An option gives one the right (but not the obligation) to buy or sell an asset, such as a stock, at a given price at a certain time in the future. They are actively traded in the financial markets and commonly priced using the Black and Scholes (1973) formula. The only unknown parameter in the Black-Scholes (BS) formula is the standard deviation of the underlying asset; this is known as the implied volatility (IV). However, whilst the BS framework assumes a constant volatility, Campa and Chang (1995) show that the IV varies as a function of time to expiration. Furthermore Canina and Figlewski (1993), Rubinstein (1994) and Foresi and Wu (2005) show that IV is asymmetric across option strikes, exhibiting a 'smile' or 'smirk'; this is also referred to as the skew. The dependence of IV upon both option strike and expiry leads to the construction of an implied volatility surface (IVS); this has also be shown to evolve over time (Cont and da Fonseca, 2002). It is these dynamics which are of particular interest for this thesis because there is a documented negative correlation between stock returns and the at-the-money (ATM) IV, an Implied Leverage Effect. This has been studied by such authors as Ciliberti et al. (2009) who have examined the relationship between the Leverage Effect and the Implied Leverage Effect.

In this chapter I present a review of the recent research on IV. Section 7.1 describes

several of the empirical findings on IV and Section 7.2 outlines several models that attempt to capture these properties. Section 7.3 presents several of the proposed explanations for the Implied Leverage Effect and Section 7.4 is a summary. For a detailed examination of the models, I refer the reader to Lee (2005).

7.1 **Properties of Implied Volatility**

It is important to understand the stylised facts of IV in order to develop better models and for risk management and trading purposes. As mentioned previously it has been widely documented that the implied volatility varies systematically with both strike and time to expiry. A great deal of research has also been conducted on the relationship between implied and realised volatility and on predicting implied volatility. Harvey and Whaley (1992) use regressions of the changes in implied volatility on information variables that include day-of-the-week dummy variables, lagged implied volatilities, interest rate measures and the lagged index return. They conclude that one-day-ahead volatility forecasts are statistically precise but do not help with devising profitable trading strategies once transaction costs are taken into account. This is consistent with the findings of Goncalves and Guidolin (2006). However Goyal and Saretto (2007) and Noh et al. (1994) were able to generate economically significant patterns even after transaction costs. The latter used a GARCH(1,1) model applied to daily changes in weighted IV. However they used a generalised least squares procedure (Day and Lewis, 1988) to compress the entire daily IVS into a single volume-weighted volatility index which Goncalves and Guidolin (2006) argue make their results less reliable.

The traditional diffusion models specify the dynamics of the spot price and its instantaneous volatility. An alternative approach is to model the dynamics of the implied volatility itself. Skiadopoulos et al. (2000) use PCA analysis, on individual smiles and the whole volatility surface, to show that 60% of the variance in the IVS can be explained by two factors: a parallel shift and a Z-shaped shift. Whilst Kamal and Derman (1997) find that three components explain 95% of the variance of the volatility surface: the level of volatility, the term structure and the skew. The first to directly model the dynamics of the IVS were Cont and da Fonseca (2002). In a study of the dynamical features of the IVS' of the S&P500 and FTSE, they show that the IVS can be described as a randomly fluctuating surface driven by a small number of factors. They also study the shapes of these factors and their dynamics. Finally, they model the IVS as a stationary random field with a covariance structure matching the empirical observations.

In explaining the dynamics of the 'smile', Daglish et al. (2007) documents two 'rules of thumb' that are commonly used in the market practice: 'Sticky-Strike' and 'Sticky-Delta/Sticky-Moneyness'. 'Sticky-Strike' assumes that the implied volatility of an option is only a function of the strike and does not depend upon the price of the underlying. Whereas 'Sticky-Moneyness' assumes that the 'smile' moves with the underlying, so the IV of a given 'moneyness' does not change. Ciliberti et al. (2009) find that empirically the 'smile' dynamics are somewhere between these two states.

7.2 Models

Local Volatility Models

The most common method for accommodating the 'smile' is to use a local volatility (LV) model such as those developed by Dupire (1994) and Derman et al. (1996). These rely on relaxing the BS assumption of constant volatility by allowing the volatility to be a function of both time and spot price; known as 'local volatility'. However, Hagan et al. (2002) finds that the dynamics of the 'smile' predicted by LV models are opposite to that observed in the market. LV models predict that when the price of the underlying asset decreases the smile shifts to higher prices and vice versa. Whilst Dumas et al. (1998) examine the predictive and hedging performance of LV models and find they are no better than an ad hoc procedure that merely smooths the IV across strike and expiry. Ciliberti et al. (2009) state that "from a more fundamental point of view, local volatility models cannot possibly represent a plausible dynamics for the underlying".

Models with Realistic Stock Dynamics

Ciliberti et al. (2009) favour models with more realistic dynamics and mention the Jump Diffusion and Levy Processes Models (Cont and Tankov, 2003), GARCH and Stochastic Volatility Models (Gatheral (2006) and Henry-Labordere (2008)) and Multifractal Models (Muzy et al. (2000) and Borland et al. (2005)). Bakshi et al. (1997) advocate using stochastic volatility models with jumps for valuing S&P500 options but Bates (2000) shows that they require unrealistic assumptions for parameter estimates.

These models predict a certain term structure for the skewness and kurtosis of the return distribution over different time scales. Having a realistic model of the stock dynamics is useful because they predict certain term structures for the skewness and kurtosis of the return distribution which can be used to estimate the smile of near-the-money options at different maturities Hagan et al. (2002), Backus et al. (1997) and Potters et al. (1998) and Bouchaud and Potters (2004). For example, iid random variables predict the skewness to decay as $T^{-1/2}$ whilst the kurtosis decays as T^{-1} . However, Ciliberti et al. (2009) state that the Leverage Effect leads to a richer term structure for skewness whilst volatility clustering leads to a non-trivial term structure for kurtosis.

Equilibrium Models

Guidolin and Timmermann (2003) show that many of the empirical biases of the BS option pricing model may be explained by Bayesian learning effects. The underlying asset price process is determined by embedding the learning mechanism in a general equilibrium model. In the model, the dividend news is allowed to evolve on a binomial lattice with unknown transition probabilities that are recursively updated using Bayes' rule. From this model they derive closed-form pricing formulas for European options and find that learning generates asymmetric skews in the IVS and systematic patterns in the term structure of option prices.
7.3 What drives the Implied Leverage Effect?

A significant factor in explaining the Implied Leverage Effect is likely to be the stock/index level Leverage Effect. Ciliberti et al. (2009) find that the stock/index level Leverage Effect gives rise to non-trivial smile dynamics and naturally explains the anomalous dependence of the skew as a function of option expiry. They also find that option markets overestimate the Implied Leverage Effect by a large factor, particularly for long dated options. This overestimation of the Implied Leverage Effect may be due to behavioural biases as suggested by Hibbert et al. (2008) and Jackwerth (2000) who recovers risk aversion functions from S&P500 options. As we saw previously, behavioural explanations are dependent upon trading behaviour and this has been studied in US stock options and S&P500 index options by Bollen and Whaley (2004). They find that changes in IV are directly related to net buying pressure from public order flow. In indices there is strong demand for put options and with no natural sellers, market makers facilitate supply and push up option prices (IV). Whereas, option writing strategies force down call option prices (IV). They also find that index options are dominated by put option trading whilst stock options are dominated by call option trading. They suggest this is why the implied volatility skew is steeper in indices than stocks.

Hibbert et al. (2008) study the Implied Leverage Effect at both the daily and intraday level and find that neither the 'Leverage Hypothesis' nor the 'Volatility Feedback Hypothesis' were able to adequately explain their results. Whilst Pena et al. (1999) use regressions and linear and non-linear Granger causality to examine 'smile' of IBEX option prices at the intraday frequency. They find that transaction costs, variables related to uncertainty about the return of the underlying asset and relative market momentum are key determinants in explaining the skew. Furthermore, they find non-linear causality effects in the dynamic interrelations between these variables and the 'smile'. In an examination of FTSE, NIKKEI and S&P500 options, Gemmill and Kamiyama (1997) find that changes in the IV in a specific market are driven by the previous period changes of IV in another market. This lagged spillover effect is also documented by Konstantinidi et al. (2008).

7.4 Summary

This chapter has outlined some of the recent research on IV for stock and index options. Importantly for this thesis, the explanation for the manifestation of Implied Leverage Effect seems reasonable, being a combination of trading activity and the stock/index level Leverage Effect. However, there remain a number of open problems. Firstly, it has proven very difficult to develop a model that exhibits realistic stock dynamics and precludes arbitrage opportunities. Secondly, there appears to be scope for more sophisticated models of the dynamics of the IVS. This seems an achievable goal since Kamal and Derman (1997) found that just three components explain 95% of the variance of the volatility surface. Hence, in the next chapter I will attempt to develop a statistical model of the IVS.

Chapter 8

Modelling Implied Volatilities using Multivariate q-Gaussian Distributions

8.1 Introduction

Cont and da Fonseca (2002) were the first to directly model the dynamics of the IVS. They modelled the IVS as a randomly fluctuating surface driven by a small number of factors. It is in this vein that I develop a descriptive statistical model of the IVS using a multivariate q-Gaussian distribution. The q-Gaussian distribution, an extension of the Gaussian distribution, naturally arises from models with multiplicative noise such as the financial markets. It has also been shown to be very useful in modelling heavy tailed data such as stock returns. I show that stock returns and IV's are well modelled by q-Gaussian distributions and identify the Implied Leverage Effect. Furthermore, I show the model may be used in practice for risk management by calculating the most probable IV changes for given stock returns. I also estimate the probability of these stock returns within a framework that is common to market practitioners but naturally allows for 'fat tails'. This could be very useful in stress testing for both vanilla and exotic derivatives.

The chapter begins by briefly discussing q-Gaussian distributions (Section 8.2) before detailing the data and calibrating the model (Section 8.3). It then proceeds to fit the q-Gaussian distribution to empirical data for S&P500 stocks (Section 8.4.1) and identifies

the Implied Leverage Effect (Section 8.4.2). Section 8.4.3 then discusses how the model can be used for stress-testing before Section 8.4.4 shows how to generate virtual data for scenario testing and option pricing using a simple Markov Chain or Auto-regressive process. Finally conclusions are presented in Section 8.5.

This research has been conducted in collaboration with Daniel Sprague. The exact responsibilities are outlined in the Declaration at the beginning of the thesis.

8.2 q-Gaussian Distributions

Tsallis (1994) developed a generalisation of Boltzmann-Gibbs thermostatistics based on the scaling properties of multifractals. He introduced a new entropy measure (or entropy functional) which was non-additive (for certain parametrisations), unlike the Boltzmann-Gibbs entropy. It is sometimes known as the Tsallis entropy. Non-extensivity arises in these systems because the entropy is not proportional to the number of elements; this is required for a system to be extensive. The financial markets are an example of such a system due to the strong correlations between assets. Financial asset returns are also characterised by intermittency and multifractal scaling for which Ramos et al. (2001) has shown the q parameter, which defines Tsallis entropy, is an objective measure.

Maximising the Tsallis entropy, under certain constraints, leads to the q-Gaussian distribution which can be considered as an extension or generalisation of the Gaussian distribution. It is characterised by it's heavy-tails (for certain parametrisations). The univariate q-Gaussian distribution is given by:

$$p(x|q,\beta) = \frac{\left(1 + \beta (q-1)x^2\right)^{\frac{1}{1-q}}}{Z_q}$$
(8.1)

where q and β are parameters to be estimated and x is the data. The normalisation factor, Z_q , is given by:

$$Z_q = \left(\beta \left(q-1\right)\right)^{\frac{1}{2}} \frac{\Gamma\left(\frac{1}{2}\right) \Gamma\left(\frac{1}{1-q}-\frac{1}{2}\right)}{\Gamma\left(\frac{1}{q-1}\right)}$$
(8.2)

This distribution is defined for q < 3 and $\beta > 0$, and it has finite variance for $q \in (-\infty, \frac{5}{3})$. For q > 1 it is 'heavy tailed' and approaches a Gaussian distribution as $q \rightarrow 1$. Figure 8.1 shows the probability density function for a q-Gaussian distribution with varying values for q and β . It shows that increasing q increases the 'heavy tailedness' of the distribution whilst increasing β appears to decrease the width of the distribution.

Some may ask why use the q-Gaussian distribution rather than another heavy tailed distribution? For example, Mandlebrot used the scaling property and the observation of power-law tails in a time series of cotton prices to propose a stable Levy distribution (with infinite variance) for financial returns. However Osorio et al. (2004) state that "Recent studies of data sets of high-frequency individual stock returns (Gopikrishnan et al. (2000) and Pleru et al. (1999)) suggest both that their variances are finite and the exponent in the power-law tail falls outside the stable Levy interval. Which opens the door to the consideration of other models to describe these distributions". The q-Gaussian distribution has been applied recently with remarkable success to describe returns of exchange rates Ramos et al. (2001) and high frequency asset returns Drozdz et al. (2007); in both the central regime and the power-law tails. Sato (2012) also computed q-Gaussian VaR measures for unconditional distributions of Japanese stocks and confirmed that returns are well fitted by q-Gaussian distributions using the Kolmogorov-Smirnov test.



Figure 8.1: Univariate q-Gaussian Distribution

This shows the Gaussian (dashed black line) and q-Gaussian distributions for a range of q and β . The q-Gaussian distribution is clearly more 'heavy tailed' than a Gaussian distribution (which is a special case of a q-Gaussian where q = 1).

Fitting q-Gaussian Distributions

In order to fit a multivariate q-Gaussian distribution one uses maximum likelihood estimation (MLE) which in low dimensions may be calculated using numerical integration. Unfortunately, estimating the normalisation factor using numerical integration becomes computationally infeasible in higher dimensions. To overcome this problem, one uses Markov-Chain Monte Carlo (*MCMC*) to sample from a target distribution, which is proportional to the desired distribution and then constructs a normalised probability density function.

Here I use the standard Metropolis-Hastings algorithm where the target distribution is

a multivariate q-Gaussian distribution:

$$p(\mathbf{x}|q,\beta) \propto \left(1 + \beta(q-1)\mathbf{x}'\Sigma^{-1}\mathbf{x}\right)^{\frac{1}{1-q}}$$
(8.3)

where q and β are parameters, Σ is the covariance matrix and **x** is the data set. The covariance matrix, Σ , is estimated from the stock data and I use a multivariate Gaussian for the proposal distribution.

Unfortunately, since one does not know the normalisation factor, one must use an approach known as Monte Carlo Maximum Likelihood Estimation (Geyer, 1991) to estimate the model parameters.

For a family of probability densities $\{f_{\theta}\}$ with respect to some measure, μ , where the densities are known only up to a normalising constant:

$$f_{\theta}(x) = \frac{1}{z(\theta)} h_{\theta}(x) \tag{8.4}$$

where h_{θ} is a known function for each θ but nothing is known about z except that:

$$z(\theta) = \int h_{\theta}(x) d\mu(x)$$
(8.5)

the integral being analytically intractable.

One can avoid calculating z by estimating the log likelihood ratio. To do this one samples a Markov Chain from any ϕ in parameter space and compare this to our target parameters, θ .

$$L(\theta) = log\left(\frac{h_{\theta}(\mathbf{x})}{h_{\phi}(\mathbf{x})}\right) - nlog\left(\frac{z(\theta)}{z(\phi)}\right)$$
(8.6)

which can be rewritten as:

$$L(\boldsymbol{\theta}) = \log\left(\frac{h_{\boldsymbol{\theta}}(\mathbf{x})}{h_{\boldsymbol{\phi}}(\mathbf{x})}\right) - n\log\left(E\left(\frac{h_{\boldsymbol{\theta}}(X_i)}{h_{\boldsymbol{\phi}}(X_i)}\right)\right)$$
(8.7)

where **x** is the data and X_i is a sample from the Markov Chain. One then minimises the negative log likelihood to estimate the model parameters.

8.3 Methodology

8.3.1 The Data

The data consists of daily prices and At-The-Money (ATM) implied volatilities for 3, 6, 12 and 18 months for S&P500 stocks during the period 2009-2014. No adjustments have been made to the data.

Figure 8.2 shows how the ATM IV varies over time, at each expiry, for Goldman Sachs (GS) stock; there are 1,984 observations. As evidenced by previous studies, the IV clearly varies over time and appears to be negatively correlated with the stock price, as observed by the IV spike when the stock drops significantly around 800 days.

Figure 8.3 shows that stock returns and changes in implied volatility for GS are all susceptible to very large tails; which can be well modelled by a q-Gaussian distribution.

The proceeding results are calculated for Goldman Sachs (GS) shares but a range of other stocks have also been examined with similar findings (Appendix D).



Figure 8.2: Time Variation of the Goldman Sachs Stock Price and At-The-Money Implied Volatility

This shows the time variation of the GS stock price and 3, 6, 12 and 18 month at-the-money implied volatilities. It clearly shows that implied volatilities vary over time and they appear to be negatively correlated with the stock price.



Figure 8.3: Time Variation of Goldman Sachs Stock Returns and Changes in At-The-Money Implied Volatility

This shows the time variation is GS stock returns and changes in 6, 12 and 18 at-the-money implied volatilities. It is clear that the implied volatility varies over time and that it exhibits very large changes.

8.3.2 Calibrating the Model

The q-Gaussian distribution is fitted using the MCMC algorithm in the 'PyMC' python library (Fonnesbeck et al., 2013). 100,000 samples were generated with a burn-in period of 10,000 and a thinning of 10. The burn-in period relates to generating a number of samples prior to recording the Markov Chain. The purpose of this is to reduce the impact of the chain starting in a low probability region of the distribution. The thinning relates to only sampling the Markov Chain every given number of observations in order to reduce auto-correlation. Both of these methods assist in reducing the convergence time of the Markov Chain. The reference values were $\beta_{\phi} = 1.50$ and $q_{\phi} = 1.30$. These were chosen to generate a more heavy tailed distribution than the data. This is important to ensure that we sample from across the entire distribution. Calibration of these parameters may often be necessary to achieve optimal convergence of the Markov Chain.

Figure 8.4 shows the trace and auto-correlation of the MCMC. It shows that the samples span a wide range and are not auto-correlated.



Figure 8.4: MCMC Calibration for Multivariate q-Gaussian

This shows the samples generated from the MCMC for GS. This is generated from a multivariate q-Gaussian distribution. (Left) the trace for the samples shows that the MCMC is sampling from a broad range. (Right) the auto-correlation for the samples shows that they are not highly autocorrelated.

8.4 Results

8.4.1 Fitting the q-Gaussian Distribution

The fitted multivariate q-Gaussian distribution has parameters: $\beta_{\theta} = 4.02$ and $q_{\theta} = 1.25$. Figure 8.5 shows the fit of the marginal distributions, with the Gaussian distribution shown for reference. Qualitatively, the results show that the q-Gaussian distribution fits the data far better than the Gaussian distribution for both returns and changes in implied volatility. The q-Gaussian appears to very closely match the actual data in the tails of the distribution. For fits of other US stocks please see Appendix D.

Figure 8.6 (Left) shows the joint distribution between returns and changes in 3 month implied volatility. The Implied Leverage Effect is clearly identifiable by the negative

correlation between the returns and the changes in implied volatility. Figure 8.6 (Right) shows the joint distribution between changes in 6 and 12 month implied volatility. There is a strong positive correlation between the implied volatilities. Both of these distributions exhibit heavy tails which are more suited to a q-Gaussian distribution.



Figure 8.5: Univariate q-Gaussian Fits of Returns and Implied Volatility Changes for Goldman Sachs

This shows the fitted marginal distributions for returns and 6, 12 and 18 month implied volatility for GS. It is clear that the q-Gaussian distribution (blue lines) fits the data far better than the Gaussian distribution. Qualitatively it also appears to fit the tails of the distribution very well.



Figure 8.6: Bivariate q-Gaussian Fits of Returns and Implied Volatility Changes for Goldman Sachs

This shows the fitted bivariate distributions for GS. (Left) this shows the joint distribution between returns and changes in 3 month implied volatility. The Leverage Effect is identifiable by the negative correlation between the returns and implied volatility. (Right) this shows the joint distribution between changes in 6 and 12 month implied volatility. It is clear that the changes in implied volatility are strongly positively correlated. Qualitatively the q-Gaussian distributions appear to fit the data reasonably well and are certainly able to capture the tails of the distribution better than a Gaussian distribution.



Figure 8.7: Fitted q-Gaussian Parameters for S&P500 Stocks

This shows the fitted parameters, (q, β) , for multivariate q-Gaussian distributions for stock returns and 3, 6, 12 and 18 month at-the-money implied volatility for S&P500 stocks. The parameters appear to be strongly positively correlated.

Figure 8.7 shows the high positive correlation between fitted values of β and q for all of the S&P 500 stocks. Averaged over all the stocks, the fitted values are $\bar{\beta} = 3.27 \pm 1.77$ and $\bar{q} = 1.24 \pm 0.04$.

As well as appearing to be qualitatively a better model, the q-Gaussian distribution is also quantitatively a better model for returns and changes in IV as shown by the Akaike Information Criterion (AIC). This measure allows a comparison between the quality of models with different numbers of parameters; a lower AIC value indicates a better model. The AIC difference between two models is calculated using:

$$\Delta AIC = 2(k_1 - k_2) - 2(\ln \mathscr{L}_1 - \ln \mathscr{L}_2)$$
(8.8)

where k_i and \mathcal{L}_i are the number of parameters and the likelihood of model *i* respectively. Please see Appendix D for further discussion of the AIC calculation for these models.

Figure 8.8 shows a histogram of the AIC difference for S&P500 stocks. The AIC value for each stock is significantly below 0, which indicates that the q-Gaussian distribution is a far better model for the empirical data than the Gaussian distribution.



Figure 8.8: Histogram of the Akaike Information Criterion for S&P500 stocks

This shows a histogram for the AIC difference between a Gaussian and a q-Gaussian for S&P500 stocks. All of the stocks exhibit a negative AIC difference which means that they are better fitted by a q-Gaussian than a Gaussian distribution. We can examine the variability of the q-Gaussian parameters over time by fitting the distribution to individual time windows. Figure 8.9 shows q and β for the multivariate q-Gaussian distribution fitted to GS stock returns and changes in 3, 6, 12 and 18 month IV for 50% overlapping time windows of 198 days. It shows that whilst q and β vary over time, they appear to do so in phase and the actual range for the q values is reasonably small given that q = 1 is a Gaussian distribution and q = 1.67 is the upper limit for a finite second moment.



Figure 8.9: Time variation of q and β for Goldman Sachs Stock

This shows the time variation of the q-Gaussian distribution parameters when fitted to 50% overlapping time windows of 198 days for the GS stock. It shows that \beta and q move in phase and the actual range of the parameters is actually quite small.

8.4.2 The Implied Leverage Effect

Figure 8.10 shows the most probable change in implied volatility conditioned on stock returns. The Implied Leverage Effect is clearly identifiable at each expiry by the negative correlation between returns and changes in implied volatility. It is also clear that the Implied Leverage Effect decreases with expiry. The decay of the Implied Leverage Effect with expiry is also shown in Figure 8.11. Here I have calculated the instantaneous skew at each expiry. This is the difference in implied volatility just above and below 0% returns. It is clear that the effect decays far more slowly than $T^{-1/2}$ which is predicted by the Black-Scholes option pricing framework.



Figure 8.10: Implied Leverage Effect for Goldman Sachs Stock

This shows the most likely change in implied volatility for a given stock return for GS. The Leverage Effect is identifiable by the negative correlation between returns and changes in implied volatility. The probability intervals for the returns as calculated by the q-Gaussian are overlayed to show the probability of a given change in implied volatility.



Figure 8.11: Decay of the Instantaneous Skew for Goldman Sachs Stock

This shows the instantaneous skew - the difference in implied volatility above and below the at-the-money strike - for GS (blue line). The results show that the instantaneous skew decays far more slowly than $1/\sqrt{t}$ as predicted in the Black-Scholes framework.

8.4.3 Stressing the Implied Volatility Surface

It is common for market practitioners to stress their derivative portfolios in order to calculate their 'Profit and Loss' and 'Greek' exposures - these represent the sensitivity of the price of a derivative to a change in underlying parameters - at different market levels. However, this can be a challenging task because the changes in the ATM IV are not constant across expires and are dependent upon the stock return itself. One can use the implied volatility dynamics in Figure 8.10 to improve on current methodologies, which tend to only consider a parallel or weighted shift of the IVS, to account for these factors. I have also included probability intervals that have been calculated in terms of standard deviations. Although standard deviations are not associated with q-Gaussian distributions, they are common in the financial markets and this makes the results comparable to standard derivative frameworks. A market practitioner can now reprice their derivative portfolio, at each market level, using these changes in IV to ascertain their 'Profit and Loss' and 'Greek' exposures. Some may argue that this does not take into consideration changes in the 'Skew/Smile' however it is very simple to add additional IV's with different 'moneyness' to this model and imply the changes in the 'Skew/Smile'.

8.4.4 Virtual Data Generation

Figure 8.12 shows that neither the GS stock returns nor the IV's are autocorrelated. In this case it is straightforward to generate virtual data via a simple Markov process:

$$\mathbf{X}_{t+1} = \mathbf{X}_t + \boldsymbol{\varepsilon} \tag{8.9}$$

where \mathbf{X}_t is a vector representing stock price and the ATM IV's at time, *t*. ε is a vector of i.i.d random variables generated from the multivariate q-Gaussian distribution. Figure 8.13 shows that qualitatively, the virtual data displays many of the same properties observed previously in the actual stock data.



Figure 8.12: Goldman Sachs Autocorrelations

Autocorrelation of Goldman Sachs returns and changes in at-the-money implied volatility for 3, 6 and 18 month expires.



Figure 8.13: Virtual Data Generation

Time variation of the virtually generated stock price and at-the-money implied volatility for 3, 6, 12 and 18 month expires using a multivariate q-Gaussian distribution with $\beta_{\theta} = 4.02$ and $q_{\theta} = 1.25$. The thinning is set to 500 to remove any autocorrelation in the MCMC.

If there are auto-correlations or cross-correlations we can incorporate them by using a vector autoregressive (VAR) framework as follows:

$$\mathbf{X}_{t+1} = \mathbf{X}_t + \sum_{j=0}^{L} \mathbf{X}_{t-j} \mathbf{B}_j + \varepsilon_t$$
(8.10)

where \mathbf{B}_j is the lag correlation matrix between t + 1 and t - j, and $\varepsilon_t \sim QG(\mathbf{0}, \Sigma, \beta, q)$ is the q-Gaussian model for the noise. By choosing the lag correlation matrix we can reproduce the observed autocorrelation and cross-correlation for any lag.

8.5 Conclusions

In this chapter I developed a descriptive statistical model of stock returns and implied volatility using a multivariate q-Gaussian distribution. Although the q-Gaussian distribution has previously been shown to fit stock returns well, this is the first time it has been used to model implied volatilities. Using empirical data for S&P500 stocks, I showed that the multivariate q-Gaussian distribution fits the stock data better than an equivalent multivariate Gaussian distribution. I then identified the Implied Leverage Effect and showed how the model could be used to estimate the probability of implied volatility changes. I also showed how this could be used for stress-testing in order to estimate portfolio 'Profit and Loss' and 'Greek' exposures. This was developed within a standard framework that could be utilised by market practitioners. In addition, I have shown how the multivariate q-Gaussian distribution could be used to generate virtual data for scenario testing and option pricing using a simple Markov Chain or Auto-regressive process. Although this analysis has only considered changes in the at-the-money implied volatility it can easily be extended to options with different 'moneyness' to account for skew and convexity effects in the implied volatility surface.

Part V

Conclusions

This thesis presented original research on the Leverage Effect in stocks, stock indices and stock options. The Leverage Effect refers to the negative correlation between an asset's return and its volatility; first documented by Black (1976). The research validated previous findings in identifying the Leverage Effect in stocks and stock indices. Contrary to the findings of Bouchaud et al. (2001), I showed that one does not need to consider the entire cross-correlation function to characterise the Leverage Effect since the persistence is actually caused by auto-information in the volatility. Furthermore, by conducting the first information theoretic analysis of the Leverage Effect, I showed that it can in fact be misleading to imply causation by simply examining cross-correlations. I also provided the first comprehensive analysis of the time varying nature of the Leverage Effect. This showed that the index level Leverage Effect was dominated by short volatile periods. This is why some authors have documented a far larger effect at the index level whilst others have questioned its statistical significance.

In ascertaining the cause of the stock level Leverage Effect I found that 50% of the information transfer between returns and volatility could be attributed to trading volumes and 42% due to an index level feedback effect. The importance of trading volumes supports the findings of Avramov et al. (2006) although they stated that the Leverage Effect could be fully explained by trading volumes. The discrepancy could arise because they only consider linear correlations. These findings were corroborated by the time variation analysis. This showed that prior to the 2008 financial crisis, the return-volume correlation led the Leverage Effect and the stock level Leverage Effect led the index level Leverage Effect. However, during and after the crisis the stock and index level Leverage Effects began to move in phase. It also showed that there was a larger information transfer from the index to the stocks during and after the crisis. This is consistent with the proposition that generally the stock Leverage Effect is driven by trading volumes but in volatile periods it is driven by systemic effects.

At the index level, a number of authors, such as Reigneron et al. (2011), have argued that the Leverage Effect is generated by stock correlation asymmetry. This posits that when the market declines a larger proportion of the stocks also decline than rise when the market rises. However, Gallant et al. (1992) claimed it could be explained by trading volumes. An information theoretic analysis showed that both of these factors played an important role with trading volumes accounting for 62% of the mutual information from returns to volatility when also conditioned on correlation asymmetry. The time variation analysis showed that prior to the 2008 financial crisis the stock level Leverage Effect led the index level Leverage Effect and both led correlation asymmetry. However, they all began move in phase during the crisis and subsequently. Therefore the correlation asymmetry could well be the mechanism by which the index level effect is amplified. This is consistent with behavioural models such as that proposed by Ahlgren et al. (2007) which they show can faithfully reproduce the index level Leverage Effect.

In order to understand how trading behaviour affects trading volumes and in turn the Leverage Effect, I utilised the model of Barberis and Xiong (2009). This allowed me to calculate the optimal trading strategy for an investor with Prospect Theory preferences. It was found that return-volume correlation is governed by the optimal trading strategy and this in turn is governed by expected returns, the standard deviation of returns and investor preferences. I validated the model predictions with a comprehensive empirical investigation of a broad range of international stocks and the trades of institutional investors in S&P500 stocks. This is one of the first pieces of research to directly link observed stylised properties with behavioural mechanisms. Prior to the 2008 financial crisis, it was found that the return-volume correlation led the Sharpe Ratio which is consistent with investors trading on the basis of expected returns and standard deviations. However, after the crisis the Sharpe Ratio led which could indicate that investors are subject to hindsight bias where they are projecting past returns into the future. This could indicate that the crisis had made the investors 'fearful' and given them negative expectations of future returns.

Finally, I developed a descriptive statistical model of stock returns and implied volatility using a multivariate q-Gaussian distribution. The q-Gaussian distribution is an extension of the Gaussian distribution that allows for better modelling of heavy tailed data. Although the q-Gaussian distribution has previously been shown to fit stock returns well, this is the first time it has been used to model implied volatilities. Using empirical data for S&P500 stocks, I showed that the multivariate q-Gaussian distribution fits the stock data better than an equivalent multivariate Gaussian distribution. I then identified the Implied Leverage Effect and showed how the model could be used to estimate the probability of implied volatility changes. I showed how this could be used for stress-testing in order to estimate portfolio 'Profit and Loss' and 'Greek' exposures. This was developed within a standard framework that could be utilised by market practitioners. I have also shown how the multivariate q-Gaussian distribution could be used to generate virtual data for scenario testing and option pricing using a simple Markov Chain or Auto-regressive process.

One of the most interesting findings of this research is the direct link between investor preferences and certain stylised properties of financial asset returns. Further investigation is required in this area to understand if investors actually behave in this manner and if so what other stylised properties are affected. This has important implications not only to investing and risk management but also to the stability of the financial system. Hens and Steude (2009) use an experimental stock market to examine the Leverage Effect. These are useful mechanisms for investigating behavioural traits as subjects behaviours can be analysed under laboratory conditions. Unfortunately, most of the experimental trading platforms - that I have seen - are not particularly realistic and do not allow participants to make informed decisions. To this end I have developed my own trading platform which has been trialled and is now ready for implementation. In association with the Psychology Department at the University of Warwick, I have constructed an experiment to examine the affect of trading behaviour on return-volume correlation. We hope that in the future, when we have sufficient funding and time, we will be able to conduct this experiment and shed further light on the impact of trading behaviour.

Bibliography

- Abdellaoui, M., Bleichrodt, H., Kammoun, H., Oct. 2011. Do Financial Professionals Behave According to Prospect Theory? An Experimental Study. Theory and Decision 74 (3), 411–429.
- Ahlgren, P., Jensen, M., Simonsen, I., Donangelo, R., Sneppen, K., 2007. Frustration Driven Stock Market Dynamics: Leverage Effect and Asymmetry. Physica A: Statistical Mechanics and its Applications 383 (1), 1–4.
- Allez, R., Bouchaud, J.-P., 2011. Individual and Collective Stock Dynamics: intra-day Seasonalities. New Journal of Physics 13 (2), 025010.
- Aloui, C., 2007. Price and Volatility Spillovers between Exchange Rates and Stock Indexes for the Pre-and Post-Euro Period. Quantitative Finance 7 (6), 669–685.
- Amblard, P., Michel, O., Dec. 2012. The Relation between Granger Causality and Directed Information Theory: A Review. Entropy 15 (1), 113–143.
- Amihud, Y., Mendelson, H., 1986. Asset Pricing and the Bid-Ask Spread. Journal of Financial Economics 17 (2), 223–249.
- Ang, A., Chen, J., 2002. Asymmetric Correlations of Equity Portfolios. Journal of Financial Economics 63 (3), 443–494.
- Ang, A., Chen, J., Xing, Y., 2006. Downside Risk. Review of Financial Studies 19 (4), 1191–1239.
- Antoniewicz, R. L., 1993. Relative Volume and Subsequent Stock Price Movements. Board of Governors, Federal Reserve System, Washington, DC.
- Atkins, A., Dyl, E., 1997. Transactions Costs and Holding Periods for Common Stocks. The Journal of Finance 52 (1), 309–325.
- Avramov, D., Chordia, T., Goyal, A., Mar. 2006. The Impact of Trades on Daily Volatility. Review of Financial Studies 19 (4), 1241–1277.

- Aydemir, A., Gallmeyer, M., Hollifield, B., 2006. Financial Leverage Does Not Cause the Leverage Effect. AFA 2007 Chicago Meetings.
- Backus, D., Foresi, S., Li, K., Wu, L., 1997. Accounting for Biases in Black-Scholes. CRIF Working Paper Series 30.
- Baek, E., Brock, W., 1992. A General Test for Nonlinear Granger Causality: Bivariate Model. Iowa State University and University of Wisconsin at Madison Working Paper.
- Bakshi, G., Cao, C., Chen, Z., 1997. Empirical Performance of Alternative Option Pricing Models. The Journal of Finance 52 (5), 2003–2049.
- Balogh, E., Simonsen, I., Nagy, B., Zoltan, N., 2010. Persistent Collective Trend in Stock Markets. Physical Review E 82 (6), 066113.
- Barberis, N., Xiong, W., Apr. 2009. What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation. The Journal of Finance 64 (2), 751–784.
- Barclay, M., Kandel, E., Marx, L., 1998. The Effects of Transaction Costs on Stock Prices and Trading Volume. Journal of Financial Intermediation 7 (2), 130–150.
- Barnett, L., Barrett, A. B., Seth, A. K., 2009. Granger Causality and Transfer Entropy are Equivalent for Gaussian Variables. Physical Review Letters 103 (23), 238701.
- Barnett, L., Bossomaier, T., Sep. 2012. Transfer Entropy as a Log-Likelihood Ratio. Physical Review Letters 109 (13), 138105.
- Baruchi, I., Grossman, D., Volman, V., Shein, M., Hunter, J., Towle, V., Ben-Jacob, E., 2006. Functional Holography Analysis: Simplifying the Complexity of Dynamical Networks. Chaos: An Interdisciplinary Journal of Nonlinear Science 16 (1), 015112.
- Bates, S., 2000. Post-'87 Crash Fears in the S&P500 Futures Option Market. Journal of Econometrics 94 (1), 181–238.
- Baur, D., Lucey, B., 2009. Flights and Contagion An Empirical Analysis of Stock-Bond Correlations. Journal of Financial Stability 5 (4), 339–352.

- Bekaert, G., Wu, G., 2000. Asymmetric Volatility and Risk in Equity Markets. Review of Financial Studies 13 (1), 1–42.
- Black, F., 1976. Studies of Stock Market Volatility Changes. In: Proceedings of the American Statistical Association, Business and Economics Section. pp. 177 – 181.
- Black, F., Scholes, M., 1973. The Pricing of Options and Corporate Liabilities. Journal of Political Economy 81, 673–654.
- Bollen, N., Whaley, R., 2004. Does Net Buying Pressure Affect the Shape of Implied Volatility Functions? The Journal of Finance 59 (2), 711–753.
- Bollerslev, T., 1986. Generalised Autoregressive Conditional Heteroskedasticity. Journal of Econometrics 31 (3), 307–327.
- Bollerslev, T., Litvinova, J., Tauchen, G., May 2006. Leverage and Volatility Feedback Effects in High-Frequency Data. Journal of Financial Econometrics 4 (3), 353–384.
- Bollerslev, T., Sizova, N., Tauchen, G., 2012. Volatility in Equilibrium: Asymmetries and Dynamic Dependencies. Review of Finance 16, 31–80.
- Bollerslev, T., Tauchen, G., Zhou, H., Feb. 2009. Expected Stock Returns and Variance Risk Premia. Review of Financial Studies 22 (11), 4463–4492.
- Borland, L., Bouchaud, J.-P., Muzy, J.-F., Zumbach, G., 2005. The Dynamics of Financial Markets: Mandelbrot's Multifractal Cascades and Beyond. Wilmott Magazine.
- Botev, Z., Grotowski, J., Kroese, D., 2010. Kernel Density Estimation via Diffusion. The Annals of Statistics 38 (5), 2916–2957.
- Bouchaud, J. P., Matacz, A., Potters, M., 2001. Leverage Effect in Financial Markets: The Retarded Volatility Model. Physical Review Letters 87 (22), 228701.
- Bouchaud, J. P., Potters, M., 2001. More Stylized Facts of Financial Markets: Leverage Effect and Downside Correlations. Physica A: Statistical Mechanics and its Applications 299 (1), 60–70.

Bouchaud, J.-P., Potters, M., 2004. Theory of Financial Risk and Derivative Pricing.

- Breeden, D. T., 1979. An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities. Journal of Financial Economics 7 (3), 265–296.
- Brooks, C., 1998. Predicting Stock Index Volatility : Can Market Volume Help? Journal of Forecasting 17 (1), 59–80.
- Campa, J., Chang, P., 1995. Testing the Expectations Hypothesis on the Term Structure of Volatilities. The Journal of Finance 50 (2), 529–547.
- Campbell, J. Y., Grossman, S. J., Wang, J., Nov. 1993. Trading Volume and Serial Correlation in Stock Returns. The Quarterly Journal of Economics 108 (4), 905–939.
- Campbell, J. Y., Hentschel, L., 1992. No News is Good News: An Asymmetric Model of Changing Volatility in Stock Returns. Journal of Financial Economics 31 (3), 281–318.
- Canina, L., Figlewski, S., 1993. The Informational Content of Implied Volatility. Review of Financial Studies 6 (3), 659–681.
- Chan, K., Fong, W.-M., 2000a. Trade Size, Order Imbalance and the Volatility-Volume Relation. Journal of Financial Economics 57 (2), 247–273.
- Chan, K., Fong, W.-M., 2000b. Trade Size, Order Imbalance and the Volatility-Volume Relation. Journal of Financial Economics 57 (2), 247–273.
- Chen, G.-M., Firth, M., Rui, O. M., Aug. 2001. The Dynamic Relation Between Stock Returns, Trading Volume, and Volatility. The Financial Review 36 (3), 153–174.
- Cheung, Y.-W., Ng, L. K., 1992. Stock Price Dynamics and Firm Size: An Empirical Investigation. The Journal of Finance 47 (5), 1985–1997.
- Chordia, T., Swaminathan, B., 2000. Trading Volume and Cross-Autocorrelations in Stock Returns. The Journal of Finance 55 (2), 913–935.
- Chou, R., 1988. Volatility Persistence and Stock Valuations: Some Empirical Evidence using GARCH. Journal of Applied Econometrics 3 (4), 279–294.

- Christie, A., 1982. The Stochastic Behavior of Common Stock Variances: Value, Leverage and Interest Rate Effects. Journal of Financial Economics 10 (4), 407–432.
- Chuang, C.-C., Kuan, C.-M., Lin, H.-Y., Jul. 2009. Causality in Quantiles and Dynamic Stock Return-Volume Relations. Journal of Banking & Finance 33 (7), 1351–1360.
- Ciliberti, S., Bouchaud, J.-P., Potters, M., 2009. Smile Dynamics: A Theory of the Implied Leverage Effect. Wilmott Journal 1 (2), 87–94.
- Clauset, A., Shalizi, C., Newman, M., 2009. Power Law Distributions in Empirical Data. SIAM review 51 (4), 661–703.
- Cont, R., 2001. Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. Journal of Quantitative Finance 1 (2), 223–236.
- Cont, R., da Fonseca, J., 2002. Dynamics of Implied Volatility Surfaces. Quantitative Finance 2 (1), 45–60.
- Cont, R., Tankov, P., 2003. Financial Modelling with Jump Processes.
- Copeland, T. E., 1976. A Model of Asset Trading under the Assumption Of Sequential Information Arrival. The Journal of Finance 31 (4), 1149–1168.
- Copeland, T. E., 1977. A Probability Model of Asset Trading. Journal of Financial and Quantitative Analysis 12 (4), 563–578.
- Cox, J. C., Huang, C.-F., Oct. 1989. Optimal Consumption and Portfolio Policies when Asset Prices follow a Diffusion Process. Journal of Economic Theory 49 (1), 33–83.
- Crouch, R. L., 1970a. A Nonlinear Test of the Random-Walk Hypothesis. American Economic Review 60 (1), 199–202.
- Crouch, R. L., 1970b. The Volume of Transactions and Price Changes on the New York Stock Exchange. Financial Analysts Journal 26 (4), 104–109.
- Daglish, T., Hull, J., Suo, W., 2007. Volatility Surfaces: Theory, Rules of Thumb and Empirical Evidence. Quantitative Finance 7 (507).

- Daouk, H., Ng, D., 2011. Is Unlevered Firm Volatility Asymmetric? Journal of Empirical Finance 18 (4), 634–651.
- Day, T., Lewis, C., 1988. The Behavior of the Volatility Implicit in the Prices of Stock Index Options. Journal of Financial Economics 22 (1), 103–122.
- Dennis, P., Mayhew, S., Stivers, C., 2006. Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon. Journal of Financial and Quantitative Analysis 41 (2), 381–406.
- Derman, E., Kani, I., Zou, J., 1996. The Local Volatility Surface: Unlocking the Information in Index Option Prices. Financial Analysts Journal 52 (4), 25–36.
- Diks, C., Panchenko, V., 2005. A note on the Hiemstra-Jones test for Granger noncausality. Studies in Nonlinear Dynamics and Econometrics 9 (2), 1558–3708.
- Diks, C., Panchenko, V., Sep. 2006. A New Statistic and Practical Guidelines for Non-Parametric Granger Causality Testing. Journal of Economic Dynamics and Control 30 (9-10), 1647–1669.
- Donangelo, R., Jensen, M., Simonsen, I., Sneppen, K., 2006. Synchronization Model for Stock Market Asymmetry. Journal of Statistical Mechanics: Theory and Experiment (11), L11001.
- Drozdz, S., Forczek, M., Kwapie, J., Rak, R., Others, 2007. Stock Market Return Distributions: From Past to Present. Physica A: Statistical Mechanics and its Applications 383 (1), 59–64.
- Duffee, G., 1995. Stock Returns and Volatility A Firm-Level Analysis. Journal of Financial Economics 37 (3), 399–420.
- Dumas, R., Fleming, J., Whaley, R., 1998. Implied Volatility Functions: Empirical Tests. The Journal of Finance 53 (6), 2059–2106.
- Dupire, B., 1994. Pricing with a Smile. Risk Magazine 7, 18-20.

- Easley, D., O'Hara, M., 1987. Price, Trade Size and Information in Securities Markets. Journal of Financial Economics 19 (1), 69–90.
- Eisler, Z., Kertesz, J., 2004. Multifractal Model of Asset Returns with Leverage Effect. Physica A: Statistical Mechanics and its Applications 343, 603–622.
- Engle, R., 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. Econometrica: Journal of the Econometric Society 50 (4), 987–1007.
- Epps, T. W., Epps, M. L., 1976. The Stochastic Dependence of Security Price Changes and Transaction Volumes: Implications for the Mixture-of-Distributions Hypothesis. Econometrica: Journal of the Econometric Society 44 (2), 305–321.
- Erb, C., Campbell, H., Tadas, V., 1994. Forecasting International Equity Correlations. Financial Analysts Journal 50 (6), 32–45.
- Fama, E., French, K., 1993. Common Risk Factors in the Returns on Stocks and Bonds. Journal of Financial Economics 33 (1), 3–56.
- Fellner, G., Maciejovsky, B., 2007. Risk Attitude and Market Behaviour: Evidence from Experimental Asset Markets. Journal of Economic Psychology 28 (3), 338–350.
- Figlewski, S., Wang, X., 2000. Is the 'Leverage Effect' a Leverage Effect? Finance Working Paper NYU.
- Finucane, F., Alhakami, A., Slovic, P., 2000a. The Affect Heuristic in Judgements of Risks and Benefits. Journal of Behavioural Decision Making 13 (1), 1–17.
- Finucane, M., Alhakami, A., Slovic, P., Johnson, S., 2000b. The Affect Heuristic in Judgements of Risks and Benefits. Journal of Behavioural Decision Making 13 (1), 1–17.
- Fonnesbeck, C., Patil, A., Huard, D., Salvatier, J., 2013. PyMC Python Module. URL http://github.com/pymc-devs/pymc

- Foresi, S., Wu, L., 2005. Crash-o-phobia: A Domestic Fear or Worldwide Concern? The Journal of Derivatives 13 (2), 8–21.
- French, K., Roll, R., 1986. Stock Return Variances: The Arrival of Information and the Reaction of Traders. Journal of Financial Economics 17 (1), 5–16.
- French, K., Schwert, G., Stambaugh, R., 1987. Expected Stock Returns and Volatility. Journal of Financial Economics 19 (1), 3–29.
- Frenzel, S., Pompe, B., 2007. Partial Mutual Information for Coupling Analysis of Multivariate Time Series. Physical Review Letters 99 (20), 204101.
- Gallant, A., Tauchen, G., 1989. Seminonparametric Estimation of Conditionally Constrained Heterogeneous Processes: Asset Pricing Applications. Econometrica: Journal of the Econometric Society 57 (5), 1091–1120.
- Gallant, A., Tauchen, G., 1993. A Nonparametric Approach to Nonlinear Time Series Analysis: Estimation and Simulation.
- Gallant, A. R., Rossi, P., Tauchen, G., Apr. 1992. Stock Prices and Volume. Review of Financial Studies 5 (2), 199–242.
- Garas, A., Argyrakis, P., 2007. Correlation Study of the Athens Stock Exchange. Physica A: Statistical Mechanics and its Applications 380, 399–410.
- Gatheral, J., 2006. The Volatility Surface: a Practitioner's Guide.
- Gemmill, G., Kamiyama, N., 1997. International Transmission of Option Volatility and Skewness: When You're Smiling, Does the Whole World Smile. City University Business School, Working Paper.
- Gennaioli, N., Shleifer, A., 2010. What Comes to Mind. The Quarterly Journal of Economics 125 (4), 1399–1433.
- Geyer, C., 1991. Markov Chain Monte Carlo Maximum Likelihood. In: Computing Science and Statistics, Proceedings of the 23rd Symposium on the Interface. pp. 156–163.

- Ghoshdastidar, D., Dukkipati, A., Bhatnagar, S., 2012. q-Gaussian based Smoothed Functional Algorithms for Stochastic Optimization. Information Theory Proceedings, 1059– 1063.
- Giraitis, L., Leipus, R., Robinson, P., Surgailis, D., 2004. LARCH, Leverage and Long Memory. Journal of Financial Econometrics 2 (2), 177–210.
- Girard, E., Biswas, R., 2007. Trading Volume and Market Volatility: Developed versus Emerging Stock Markets. Financial Review 42 (3), 429–459.
- Glosten, L., Jaganathan, R., Runkle, D., 1993. On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. The Journal of Finance 48 (5), 1779–1801.
- Godfrey, M., Granger, C., Morgenstern, O., 1964. The Random Walk Hypothesis of Stock Market Behaviour. Kyklos 17 (1), 1–30.
- Gomez-Herrero, G., Wu, W., 2010. Assessing Coupling Dynamics from an Ensemble of Time Series. arXiv preprint arXiv:1008.0539.
- Goncalves, S., Guidolin, M., 2006. Predictable Dynamics in the S&P500 Index Options Implied Volatility Surface. The Journal of Business 79 (3), 1591–1635.
- Gopikrishnan, P., Pleru, V., Gabaix, X., Stanley, H., 2000. Scaling and Correlation in Financial Time Series. Physica A 287, 362–373.
- Goyal, A., Saretto, A., 2007. Option Returns and Volatility Mispricing. Purdue University, Working Paper.
- Graham, J., Harvey, C., 2001. Expectations of Equity Risk Premia, Volatility and Asymmetry from a Corporate Finance Perspective. National Bureau of Economic Research.
- Granger, C., Morgenstern, O., 1963. Spectral Analysis of New York Stock Market Prices. Kyklos 16 (1), 1–27.

- Granger, C. W. J., 1969. Investigating Causal Relations by Econometric Models and Cross-Spectral Methods. Econometrica: Journal of the Econometric Society 37 (3), 424–438.
- Graps, A., 1995. An Introduction to Wavelets. Computational Science & Engineering, IEEE 2 (2), 50–61.
- Guidolin, M., Timmermann, A., 2003. Option Prices Under Bayesian Learning: Implied Volatility Dynamics and Predictive Densities. Journal of Economic Dynamics and Control 27 (5), 717–769.
- Hagan, P., Kumar, D., Lesniewski, A., 2002. Managing Smile Risk. Wilmott Magazine, 84–108.
- Harris, L., 1983. The Joint Distribution of Speculative Prices and of Daily Trading Volume. Unpublished Manuscript, University of Southern California.
- Harris, L., Gurel, E., 1986. Price and Volume Effects Associated with Changes in the S&P
 500 List: New Evidence for the Existence of Price Pressures. The Journal of Finance
 41 (4), 815–829.
- Harris, M., Raviv, A., 1993. Differences of Opinion make a Horse Race. Review of Financial Studies 6 (3), 473–506.
- Harvey, C., Whaley, R., 1992. Market Volatility Prediction and the Efficiency of the S&P100 Index Option Market. Journal of Financial Economics 31 (1), 43–73.
- Hasanhodzic, J., Lo, A., Feb. 2011. Black's Leverage Effect is not due to Leverage. SSRN Electronic Journal.
- He, H., Wang, J., 1995. Differential Informational and Dynamic Behavior of Stock Trading Volume. Review of Financial Studies 8 (4), 919–972.
- Hellwig, M., 1980. On the Aggregation of Information in Competitive Markets. Journal of Economic Theory 22 (3), 477–498.

- Henry-Labordere, P., 2008. Analysis, Geometry and Modelling in Finance: Advanced Methods in Option Pricing.
- Hens, T., Steude, S. C., Jun. 2009. The Leverage Effect Without Leverage. Finance Research Letters 6 (2), 83–94.
- Heston, S., Nandi, S., 2000. A Closed-Form GARCH Option Valuation Model. Review of Financial Studies 13 (3), 585–625.
- Hibbert, A.-M., Daigler, R. T., Dupoyet, B., Oct. 2008. A Behavioral Explanation for the Negative Asymmetric Return-Volatility Relation. Journal of Banking & Finance 32 (10), 2254–2266.
- Hiemstra, C., Jones, J. D., Dec. 1994. Testing for Linear and Nonlinear Granger Causality in the Stock Price-Volume Relation. The Journal of Finance 49 (5), 1639–1664.
- Hlaváčková-Schindler, K., 2011. Equivalence of Granger Causality and Transfer Entropy : A Generalization. Applied Mathematical Sciences 5 (73), 3637–3648.
- Hlaváčková-Schindler, K., Paluš, M., 2007. Causality Detection Based on Information-Theoretic Approaches in Time Series Analysis. Physics Reports 441 (1), 1–46.
- Jackwerth, J., 2000. Recovering Risk Aversion from Option Prices and Realized Returns. Review of Financial Studies 13 (2), 433–451.
- Jennings, R. H., Starks, L., Fellingham, J., 1981. An Equilibrium Model of Asset Trading with Sequential Information Arrival. The Journal of Finance 36 (1), 143–161.
- Jones, C., Kaul, G., Lipson, M., 1994a. Transactions, Volume and Volatility. Review of Financial Studies 7 (4), 631–651.
- Jones, C., Kaul, G., Lipson, M., 1994b. Transactions, Volume and Volatility. Review of Financial Studies 7 (4), 631–651.
- Jung, W.-S., Chae, S., Yang, J.-S., Moon, H.-T., 2006. Characteristics of the Korean Stock Market Correlations. Physica A: Statistical Mechanics and its Applications 361 (1), 263–271.

- Kahneman, D., Tversky, A., 1979. Prospect Theory: An Analysis of Decision under Risk. Econometrica 47 (2), 263–292.
- Kamal, M., Derman, E., 1997. The Patterns of Change in Implied Index Volatilities. Goldman Sachs, Quantitative Strategies Notes.
- Karpoff, J. M., Mar. 1987. The Relation between Price Changes and Trading Volume: A Survey. Journal of Financial and Quantitative Analysis 22 (01), 109–126.
- Kenett, D., Shapira, Y., Madi, A., Bransburg-Zabary, S., Gur-Gershgoren, G., Ben-Jacob,E., 2011. Index Cohesive Force Analysis Reveals that the US Market became Prone toSystemic Collapses since 2002. PloS one 6 (4), e19378.
- Kenett, D., Shapira, Y., Madi, A., Bransburg-Zabary, S., Gus-Gershgoren, G., Ben-Jacob,
 E., 2010. Dynamics of Stock Market Correlations. AUCO Czech Economic Review
 4 (3), 330–341.
- Kogan, L., Sep. 2004. Asset Prices and Real Investment. Journal of Financial Economics 73 (3), 411–431.
- Konstantinidi, E., Skiadopoulos, G., Tzagkaraki, E., 2008. Can the Evolution of Implied Volatility be Forecasted? Evidence from European and US Implied Volatility Indices. Journal of Banking and Finance 32 (11), 2401–2411.
- Kullback, S., Leibler, R., 1951. On Information and Sufficiency. The Annals of Mathematical Statistics 22 (1), 79–86.
- Kwon, O., Oh, G., 2012. Asymmetric Information Flow between Market Index and Individual Stocks in Several Stock Markets. EPL (Europhysics Letters) 97 (2), 28007.
- Kwon, O., Yang, J.-S., 2008a. Information Flow between Composite Stock Index and Individual Stocks. Physica A: Statistical Mechanics and its Applications 387 (12), 2851– 2856.
- Kwon, O., Yang, J.-S., 2008b. Information Flow between Stock Indices. EPL (Europhysics Letters) 82 (6), 68003.

- Laloux, L., Cizeau, P., Bouchaud, J.-P., Potters, M., 1999. Noise Dressing of Financial Correlation Matrices. Physical Review Letters 83 (7), 1467.
- Lee, B.-S., Rui, O. M., Jan. 2002. The Dynamic Relationship between Stock Returns and Trading Volume: Domestic and Cross-Country Evidence. Journal of Banking & Finance 26 (1), 51–78.
- Lee, R., 2005. Implied Volatility: Statics, Dynamics, and Probabilistic Interpretation. Recent Advances in Applied Probability, 241–268.
- Lleras, A., Kawahara, J., Levinthal, B., 2009. Past Rejections Lead to Future Misses : Selection-related Inhibition Produces Blink-like Misses of Future (easily detectable) Events. Journal of Vision 9 (26), 1–12.
- Llorente, G., Michaely, R., Saar, G., Wang, J., Jul. 2002. Dynamic Volume-Return Relation of Individual Stocks. Review of Financial Studies 15 (4), 1005–1047.
- Ma, W.-J., Hu, C.-K., Amritkar, R., 2004. Stochastic Dynamical Model for Stock-Stock Correlations. Physical Review E 70 (2), 026101.
- Mathworks, 2014. Matlab Wavelet Families.

Michael, F., Johnson, M., ???? Financial Market Dynamics. Physica A.

- Morgan, I. G., 1976. Stock Prices and Heteroskedasticity. Journal of Business 49 (4), 496–508.
- Mulherin, J. H., Gerety, M. S., 1988. Trading Volume on the NYSE during the Twentieth Century: A Daily and Hourly Analysis. US Securities and Exchange Commission, Office of Economic Analysis.
- Muzy, J.-F., Delour, J., Bacry, E., 2000. Modelling Fluctuations of Financial Time Series: from Cascade Process to Stochastic Volatility Model. The European Physical Journal B 17 (3), 537–548.
- Nelson, D., 1991. Conditional Heteroskedasticity in Asset Returns: A New Approach. Econometrica: Journal of the Econometric Society 59 (2), 347–370.
- Neumaier, A., Schneider, T., 2001. Estimation of Parameters and Eigenmodes of Multivariate Autoregressive Models. ACM Transactions on Mathematical Software 27 (1), 27–57.
- Noh, J., 2000. Model for Correlations in Stock Markets. Physical Review E 61 (5), 5981.
- Noh, J., Engle, R., Kane, A., 1994. Forecasting Volatility and Option Prices of the S&P500 Index. The Journal of Derivatives 2 (1), 17–30.
- Odean, T., 1998. Are Investors Reluctant to Realise their Losses? The Journal of Finance 53 (5), 1775–1798.
- Osorio, R., Borland, L., Tsallis, C., 2004. Distributions of High-Frequency Stock-Market Observables. Nonextensive Entropy: Interdisciplinary Applications, 321.
- Pan, R., Sinha, S., 2007. Collective Behavior of Stock Price Movements in an Emerging Market. Physical Review E 76 (4), 046116.
- Papana, A., Kugiumtzis, D., 2009. Evaluation of Mutual Information Estimators for Time Series. International Journal of Bifurcation and Chaos 19 (12), 4197–4215.
- Pena, I., Rubio, G., Serna, G., 1999. Why do we Smile? On the Determinants of the Implied Volatility Function. Journal of Banking & Finance 23 (8), 1151–1179.
- Pindyck, R., 2008. Uncertainty in the Theory of Renewable Resource Markets. The Review of Economic Studies 51 (2), 289–303.
- Plerou, V., Parameswaran, G., Rosenow, B., Amaral, L., Guhr, T., Stanley, H. E., 2002. Random Matrix Approach to Cross Correlations in Financial Data. Physical Review E 65 (6), 066126.
- Pleru, V., Gopikrishnan, P., Amaral, L., Meyer, M., Stanley, H., 1999. Scaling of the Distributions of Price Fluctuations of Individual Companies. Physical Review E 60, 6519–6529.

- Pontiff, J., Woodgate, A., 2008. Share Issuance and Cross-Sectional Returns. The Journal of Finance 63 (2), 921–945.
- Poterba, J., Summers, L., 1986. Poterba persistence of volatility and stock market fluctuations. American Economic Review, 1142–1151.
- Potters, M., Cont, R., Bouchaud, J.-P., 1998. Financial Markets as Adaptive Ecosystems. European Physical Letters 41, 239.
- Qiu, T., Zheng, B., Ren, F., Trimper, S., 2006. Return-volatility Correlation in Financial Dynamics. Physical Review E 73 (6), 065103.
- Qiu, T., Zheng, B., Ren, F., Trimper, S., 2007. Statistical Properties of German Dax and Chinese Indices. Physica A: Statistical Mechanics and its Applications 378 (2), 387– 398.
- Ramos, F., Rodrigues Neto, C., Rosa, R., Abreu Sa, L., Bolzam, M., 2001. Generalized Thermostatistical Description of Intermittency and Nonextensivity in Turbulence and Financial Markets. Nonlinear Analysis 23, 3521–3530.
- Reigneron, P.-A., Allez, R., Bouchaud, J.-P., 2011. Principal Regression Analysis and the Index Leverage Effect. Physica A: Statistical Mechanics and its Applications 390 (17), 3026–3035.
- Rogalski, R., 1979. The Dependence of Prices and Volume. Review of Economics and Statistics 60 (2), 268–274.
- Roman, H., Porto, M., Dose, C., 2008. Skewness, Long-Time Memory, and Non-Stationarity: Application to Leverage Effect in Financial Time Series. EPL (Europhysics Letters) 84 (2), 28001.
- Rubinstein, M., 1994. Implied Binomial Trees. The Journal of Finance 49 (3), 771-818.
- Saatcioglu, K., Starks, L. T., Jun. 1998. The Stock Price-Volume Relationship in Emerging Stock Markets: the Case of Latin America. International Journal of Forecasting 14 (2), 215–225.

- Sato, A., 2010. q-Gaussian Distributions and Multiplicative Stochastic Processes for Analysis of Multiple Financial Time Series. Journal of Physics: Conference Series 201 (1), 012008.
- Sato, A., 2012. A Method to Quantify Risks of Financial Assets: An Empirical Analysis of Japanese Security Prices. Advanced Materials Research 452, 469–473.
- Schreiber, T., 2000. Measuring Information Transfer. Physical Review Letters 85 (2), 461.
- Schwert, G. W., 1989. Why Does Stock Market Volatility Change Over Time? The Journal of Finance 44 (5), 1115–1153.
- Shannon, C. E., 1948. A Note on the Concept of Entropy. Bell System Tech. J.
- Shapira, Y., Kenett, D. Y., Ben-Jacob, E., Nov. 2009. The Index cohesive effect on stock market correlations. The European Physical Journal B 72 (4), 657–669.
- Shapira, Y., Kenett, D. Y., Raviv, O., Ben-Jacob, E., 2011. Hidden temporal order unveiled in stock market volatility variance. AIP Advances 1 (2), 022127.
- Shefrin, H., 2001. Behavioural Corporate Finance. Journal of Applied Corporate Finance 14 (3), 113–126.
- Shefrin, H., 2007. Behavioral corporate finance: Decisions that create value.
- Shen, J., Zheng, B., 2009. On Return-Volatility Correlation in Financial Dynamics. EPL (Europhysics Letters) 88 (2), 28003.
- Simonsen, I., Ahlgren, P., Jensen, M., Donangelo, R., Sneppen, K., 2007. Fear and its Implications for Stock Markets. The European Physical Journal B 57 (2), 153–158.
- Simonsen, I., Jensen, M., Johansen, A., 2002. Optimal Investment Horizons. The European Physical Journal B 27 (4), 583–586.
- Siven, J., Lins, J., 2009. Temporal Structure and Gain-Loss Asymmetry for Real and Artificial Stock Indices. Physical Review E 80 (5), 057102.

- Skiadopoulos, G., Hodges, S., Clewlow, L., 2000. The Dynamics of the S&P500 Implied Volatility Surface. Review of Derivatives Research 3 (3), 263–282.
- Smirlock, M., Starks, L., 1988. An Empirical Analysis of the Stock Price-Volume Relationship. Journal of Banking & Finance 12 (1), 31–41.
- Smirnov, D., Apr. 2013. Spurious Causalities with Transfer Entropy. Physical Review E 87 (4), 042917.
- Smith, L. V., Yamagata, T., 2011. Firm Level Return-Volatility Analysis Using Dynamic Panels. Journal of Empirical Finance 18 (5), 847–867.
- Stickel, S. E., Verrecchia, R. E., Nov. 1994. Evidence that Trading Volume Sustains Stock Price Changes. Financial Analysts Journal 50 (6), 57–67.
- Struzik, R., 2001. Wavelet Methods in (Financial) Time-Series Processing. Physica A 297, 307–319.
- Talpsepp, T., Reiger, M., 2010. Explaining Asymmetric Volatility Around the World. Journal of Empirical Finance 17 (5), 938–956.
- Tauchen, G., 2011. Stochastic Volatility in General Equilibrium. The Quarterly Journal of Finance 1 (4), 707–731.
- Thompson, S. B., 2011. Simple Formulas for Standard Errors that Cluster by both Firm and Time. Journal of Financial Economics 99 (1), 1–10.
- Torrence, C., Campo, G., 1998. A Practical Guide to Wavelet Analysis. Bulletin of the American Meterological Society 79 (1), 61–78.
- Tsallis, C., 1994. What are the numbers that experiments provide? Quimica Nova 17 (6), 468–471.
- Turner, C., Startz, R., Nelson, C., 1989. A Markov model of heteroskedasticity, risk, and learning in the stock market. Journal of Financial Economics 25 (1), 3–22.

- Tversky, A., Kahneman, D., Oct. 1992. Advances in Prospect Theory: Cumulative Representation of Uncertainty. Journal of Risk and Uncertainty 5 (4), 297–323.
- Utsugi, A., Ino, K., Oshikawa, M., 2004. Random Matrix Theory Analysis of Cross Correlations in Financial Markets. Physical Review E 70 (2), 026110.
- Visaltanachoti, N., Lu, L., Luo, H., 2007. Holding Periods, Illiquidity and Disposition Effect in the Chinese Stock Markets. Applied Financial Economics 17 (15), 1265–1274.
- Wang, J., 1994. A Model of Competitive Stock Trading Volume. Journal of Political Economy 102 (1), 127–168.
- Wang, J., Yang, M., 2009. Asymmetric Volatility in the Foreign Exchange Markets. Journal of International Financial Markets, Institutions and Money 19 (4), 597–615.
- Westerfield, R., 1977. The Distribution of Common Stock Price Changes: An Application of Transactions Time and Subordinated Stochastic Models. Journal of Financial and Quantitative Analysis 12 (5), 743–765.
- Wiener, N., 1956. The Theory of Prediction. In: Modern Mathematics for Engineers.Vol. 1. McGraw-Hill: New York, NY, USA, pp. 165–183.
- Wood, R. A., McInish, T. H., Ord, J. K., 1985. An Investigation of Transactions Data for NYSE Stocks. The Journal of Finance 40 (3), 723–739.
- Wu, G., 2001. The Determinants of Asymmetric Volatility. The Review of Financial Studies 14 (3), 837–859.
- Yang, S.-Y., Doong, S.-C., 2004. Price and Volatility Spillovers between Stock Prices and Exchange Rates: Empirical Evidence from the G-7 Countries. International Journal of Business and Economics 3 (2), 139–153.
- Ying, C., 1996. Stock Market Prices and Volumes of Sales. Econometrica 34 (3), 676–685.

Appendices

Appendix A

An Information Theoretic Analysis of Returns, Volatility and Trading Volumes at the Weekly Frequency

In Chapters 2 and 5 I examined the dynamic relationships between returns, volatility and trading volumes at the daily frequency for stocks and stock indices respectively. The problem with analysing the data at the daily frequency is that the *MI* struggles to differentiate volatility and returns. This is because volatility is simply defined as the square of the daily returns. Here I repeat the analysis in Chapters 2 and 5 at the weekly frequency. This enables me to decouple returns and volatility by defining volatility as the sum of the daily squared returns over the week. Section A.1 examines the stock level effect - the return-volatility relationship is examined in Chapter 2 - whilst Chapter A.2 examines the index level effect.

A.1 Stock Results

A.1.1 Calibrating the Model



Figure A.1: Mutual Information Calibration for Stocks at the Weekly Frequency

Shows the mean *MI* for S&P500 stocks, at the weekly frequency, as a function of the number of nearest neighbours, *k*. The blue stars with dashed line is the $MI(\hat{r}_t, \hat{v}_{t+1})$, the red diamonds with solid line is the $MI(\hat{r}_t, \hat{x}_{t+1})$ and the green squares with dot-dashed line is $MI(\hat{x}_t, \hat{v}_{t+1})$. These are given with associated one standard errors. The results appear to show that the *MI* converges as the number of nearest neighbours increases and that $k \approx 75$ should be sufficient to produce stable results.

A.1.2 Persistence of Returns, Volatility and Volumes



Figure A.2: Auto-Mutual Information and Auto-Covariance Functions for Stocks at the Weekly Frequency

Shows the auto-mutual information (Top) and normalised auto-covariance (Bottom) functions. They are calculated for returns (blue stars with dashed lines), volatility (green squares with dot-dashed lines) and volumes (red diamonds with solid lines). These are given with associated one standard errors, exponential curve fits and 95% Significance Levels (black dashed lines). The *MI* is statistically significant and persistent for volatility and volumes with the *MI* for volumes significantly larger than volatility. The *MI* is only statistically significant for returns for the first 4 weeks. The auto-covariance function exhibits similar properties to the auto-mutual information function.

A.1.3 Returns and Volume



Figure A.3: Cross-Mutual Information and Cross-Covariance Functions for Stock Returns and Volume at the Weekly Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions between returns and volume. The blue squares with solid lines show the $MI(\hat{x}_t, \hat{r}_{t+\tau})$ and $Corr(x_t, r_{t+\tau})$. The green diamonds with dot-dashed lines show the $MI(\hat{r}_t, \hat{x}_{t+\tau})$ and $Corr(r_t, x_{t+\tau})$. These are given with associated one standard errors, stretched exponential curve fits and 95% Significance Levels (dashed black lines). The MI is statistically significant only for $MI(\hat{r}_t, \hat{x}_{t+\tau})$ for $\tau < 4$. This indicates returns cause volumes and not vice versa. This is consistent with the cross-covariance function which only shows structure for $Corr(r_t, x_{t+\tau})$ although this is not statistically significant.



Figure A.4: Cross-Partial Mutual Information and Cross-Transfer Entropy Functions for Stock Returns and Volume at the Weekly Frequency

(Top) Shows the partial cross-mutual information function for returns and volume, where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{x}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised returns. It is given with associated one standard errors, stretched exponential curve fits and a 95% Significance Level (dashed black line). The *PMI* is not statistically significant in either direction. (Bottom) Shows the cross-transfer entropy function for returns and volumes. The green diamonds represent $TE(\hat{r}_t \to \hat{x}_{t+\tau})$ and the blue squares represent the $TE(\hat{x}_t \to \hat{r}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \to \hat{x}_{t+\tau})$ and the dashed blue line for $TE(\hat{x}_t \to \hat{r}_{t+\tau})$. The *TE* shows that $TE(\hat{x}_t \to \hat{r}_{t+\tau}) > TE(\hat{r}_t \to \hat{x}_{t+\tau})$ for all τ but neither are not statistically significant.

A.1.4 Volume and Volatility



Figure A.5: Cross-Mutual Information and Cross-Covariance Functions for Stock Volatility and Volume at the Weekly Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for volumes and volatility. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau})$ and the $Cov(v_t, x_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau})$ and the $Cov(x_t, v_{t+\tau})$. Both graphs are given with associated one standard errors, stretched exponential curve fits and 95% Significance Levels (black dashed lines). The *MI* is statistically significant and persistent in both directions with $MI(\hat{v}_t, \hat{x}_{t+1}) > MI(\hat{x}_t, \hat{v}_{t+1})$ but approximately equal for $\tau \ge 2$. This indicates bi-directional causality between volume and volatility. This is consistent with the results of the cross-covariance function although $Cov(v_t, x_{t+\tau})$ and $Cov(x_t, v_{t+\tau})$ are only statistically significant for $\tau \le 4$.



Figure A.6: Cross-Partial Mutual Information and Cross-Transfer Entropy Functions for Stock Volatility and Volume at the Weekly Frequency

(Top) Shows the partial cross-mutual information function for volumes and volatility where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volatility. It is given with associated one standard errors, stretched exponential curve fits and a 95% Significance Level (black dashed line). The MI just statistically significant for $MI(\hat{x}_t, \hat{v}_{t+1} | \hat{v}_t)$ which indicates that volumes cause volatility and not vice versa. Since neither MI is statistically significant for $\tau > 1$, the PMI also implies that the persistence is due to auto-information.. (Bottom) Shows the cross-transfer entropy function for volumes and volatility. The green diamonds represent $TE(\hat{v}_t \to \hat{x}_{t+\tau})$ and the blue squares represent $TE(\hat{x}_t \to \hat{v}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels where the green dot-dashed line is for the $TE(\hat{v}_t \to \hat{x}_{t+\tau})$ and the dashed blue line is for the $TE(\hat{x}_t \to \hat{v}_{t+\tau})$. This shows that $TE(\hat{x}_t \to \hat{v}_{t+\tau}) > TE(\hat{v}_t \to \hat{x}_{t+\tau})$ for all τ and statistically significant which implies volumes (Granger) cause volatility and not vice versa.

A.2 Index Results

A.2.1 Calibrating the Model



Figure A.7: Mutual Information Calibration for Indices at the Weekly Frequency

Shows the mean *MI* for developed market indices, at the weekly frequency, as a function of the number of nearest neighbours, *k*. The blue stars with dashed line is the $MI(\hat{r}_t, \hat{v}_{t+1})$, the red diamonds with solid line is the $MI(\hat{r}_t, \hat{x}_{t+1})$ and the green squares with dot-dashed line is $MI(\hat{x}_t, \hat{v}_{t+1})$. These are given with associated one standard errors. The results appear to show that the *MI* converges as the number of nearest neighbours increases and that $k \approx 75$ should be sufficient to produce stable results.

A.2.2 Persistence of Returns, Volatility and Volumes



Figure A.8: Auto-Mutual Information and Auto-Covariance Functions for Indices at the Weekly Frequency

Shows the auto-mutual information (Top) and normalised auto-covariance (Bottom) functions. They are calculated for returns (blue stars with dashed lines), volatility (green squares with dot-dashed lines) and volumes (red diamonds with solid lines). These are given with associated one standard errors, exponential curve fits and 95% Significance Levels (black dashed lines). The *MI* is statistically significant and persistent for volatility and volumes with the *MI* for volumes significantly larger than volatility. The *MI* is only statistically significant for returns for the first 10 weeks. The auto-covariance function shows that the volumes are highly persistent but the volatility decays within 15 weeks whilst the returns are not statistically significant at any time horizon.

A.2.3 Returns and Volume



Figure A.9: Cross-Mutual Information and Cross-Covariance Functions for Index Returns and Volume at the Weekly Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions between returns and volume. The blue squares with solid lines show the $MI(\hat{x}_t, \hat{r}_{t+\tau})$ and $Corr(x_t, r_{t+\tau})$. The green diamonds with dot-dashed lines show the $MI(\hat{r}_t, \hat{x}_{t+\tau})$ and $Corr(r_t, x_{t+\tau})$. These are given with associated one standard errors, exponential curve fits and 95% Significance Levels (dashed black lines). The *MI* is statistically significant and persistent which implies bi-directional causality. This contrasts with the cross-covariance function which is not statistically significant at any time horizon.



Figure A.10: Cross-Partial Mutual Information and Cross-Transfer Entropy Functions for Index Returns and Volume at the Weekly Frequency

(Top) Shows the partial cross-mutual information function for returns and volume, where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{x}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised returns. It is given with associated one standard errors, exponential curve fits and a 95% Significance Level (dashed black line). The *PMI* is only statistically significant for $MI(\hat{x}_t, \hat{r}_{t+1}|\hat{r}_t)$ which implies volumes (Granger) cause returns and not vice versa. Since neither MI is statistically significant for $\tau > 1$, the *PMI* also implies that the persistence is due to auto-information. (Bottom) Shows the cross-transfer entropy function for returns and volumes. The green diamonds represent $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ and the blue squares represent the $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ and the dashed blue line for $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau})$. The *TE* shows that $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau}) > TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ for all τ and statistically significant. This implies volumes (Granger) cause returns and not vice versa.

A.2.4 Volume and Volatility



Figure A.11: Cross-Mutual Information and Cross-Covariance Functions for Stock Volatility and Volume at the Weekly Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for volumes and volatility. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau})$ and the $Cov(v_t, x_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau})$ and the $Cov(x_t, v_{t+\tau})$. Both graphs are given with associated one standard errors, exponential curve fits and 95% Significance Levels (black dashed lines). The MI is statistically significant and persistent in both directions with $MI(\hat{v}_t, \hat{x}_{t+1}) \approx MI(\hat{x}_t, \hat{v}_{t+1})$. This indicates bi-directional causality between volume and volatility. This is consistent with the results of the cross-covariance function, although $Cov(v_t, x_{t+\tau})$ and $Cov(x_t, v_{t+\tau})$ are only statistically significant for around 5 weeks.



Figure A.12: Cross-Partial Mutual Information and Cross-Transfer Entropy Functions for Index Volatility and Volume at the Weekly Frequency

(Top) Shows the partial cross-mutual information function for volumes and volatility where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volatility. It is given with associated one standard errors, exponential curve fits and a 95% Significance Level (black dashed line). The MI is only statistically significant for $MI(\hat{x}_t, \hat{v}_{t+\tau} \mid Z)$ for $\tau < 3$ which indicates that volumes cause volatility and not vice versa. Since neither MI is statistically significant for $\tau > 4$, the PMI also implies that the persistence is due to auto-information. (Bottom) Shows the cross-transfer entropy function for volumes and volatility. The green diamonds represent $TE(\hat{v}_t \to \hat{x}_{t+\tau})$ and the blue squares represent $TE(\hat{x}_t \to \hat{v}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels where the green dot-dashed line is for the $TE(\hat{v}_t \to \hat{x}_{t+\tau})$ and the dashed blue line is for the $TE(\hat{x}_t \to \hat{v}_{t+\tau})$. This shows that $TE(\hat{x}_t \to \hat{v}_{t+\tau}) > TE(\hat{v}_t \to \hat{x}_{t+\tau})$ for all τ and statistically significant which implies volumes (Granger) cause volatility and not vice versa.

A.2.5 Returns and Volatility



Figure A.13: Cross-Mutual Information and Cross-Covariance Functions for Index Returns and Volatility at the Weekly Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for returns and volatility at the weekly frequency. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau})$ and the $Cov(r_t, v_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau})$ and the $Cov(v_t, r_{t+\tau})$. Both graphs are given with associated one standard errors, exponential curve fits and 95% Significance Levels (black dashed lines). The cross-mutual information function shows that $MI(\hat{r}_t, \hat{v}_{t+\tau}) > MI(\hat{v}_t, \hat{r}_{t+\tau})$ for $\tau \leq 8$ but they are both statistically significant. This implies a bi-directional information flow between returns and volatility. The cross-covariance function only shows structure for $Cov(r_t, v_{t+\tau})$ which is statistically significant for $\tau < 3$. This shows that the Leverage Effect is also observable at the weekly frequency.



Figure A.14: Cross-Partial Mutual Information and Cross-Transfer Entropy Functions for Index Returns and Volatility at the Weekly Frequency

(Top) Shows the partial cross-mutual information function for returns and volatility at the weekly frequency. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volatilities. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised returns. These are given with associated one standard errors, exponential curve fits and a 95% Significance Level (dashed black line). It shows that $MI(\hat{v}_t, \hat{r}_{t+\tau}|Z)$ for $\tau < 4$ is statistically significant as is $MI(\hat{r}_t, \hat{v}_{t+1}|Z)$. This is a stronger indication that volatility cause returns at the weekly frequency and not vice versa. (Bottom) Shows the cross-transfer entropy functions for returns and volatility at the weekly frequency. The green diamonds represent the $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the blue squares represent the $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. These are given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the dashed blue line for $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. The *TE* results support those of the *PMI* but the $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$ is now statistically significant for $\tau > 20$ weeks. This indicates that volatility (Granger) cause returns at the weekly frequency and not vice versa.

Appendix B

Information Theoretic Analysis of Returns, Volatility and Trading Volumes by Index

In Chapter 2 I examined the dynamic relationships between stock returns, volatility and trading volumes for S&P500 stocks. Here I extend that analysis to consider stocks from a broad range of global indices. The results are generally consistent with previous findings.

B.1 The Data

The data is sourced from Bloomberg and covers daily returns and volumes for a range of global stocks during the period 1980-2012. The data is not corrected for stocks that have been added/removed from the indices. However, I do not believe that this unduly affects the results. All statistics and results have been calculated at the individual stock level and then averaged to give a value for the overall index. The summary statistics are shown in Table B.1.

Country	Country Index		Mean Return	Mean Std	Mean Volume
		Stocks			
Argentina	MERVAL	10	0.021%	0.034	0.907
Australia	ASX	176	0.033%	0.033	2.343
Brazil	BOVESPA	50	0.042%	0.034	1.890
Canada	TSX	216	0.041%	0.031	0.574
China	SHANGHAI	846	0.014%	0.032	7.378
France	CAC	40	0.014%	0.032	2.316
Germany	DAX	30	0.014%	0.032	3.357
Holland	AEX	22	0.014%	0.032	3.724
Hong	HSI	43	0.016%	0.032	38.034
Kong					
India	SENSEX	28	0.017%	0.032	2.305
Italy	MIB	33	0.014%	0.032	11.419
Japan	NKY	214	0.002%	0.025	3.702
Korea	KOSPI	620	0.015%	0.037	0.340
Mexico	MEXBOL	28	0.068%	0.027	5.612
Spain	IBEX	30	0.020%	0.021	5.591
Sweden	OMX	28	0.039%	0.024	3.679
Switzerland	SMI	18	0.026%	0.022	2.858
UK	FTSE	93	0.035%	0.023	8.819
USA	NDX	87	0.066%	0.035	5.982
USA	SPX	454	0.041%	0.026	2.300

Table B.1: Summary Statistics by Index

Shows the number of stocks, the mean daily return, standard deviation and volume for each index in the data set.

B.2 Results by Index

B.2.1 Returns and Volume

Table B.2: Summary of the Stock Level Return-Volume Re-lation by Index

Index	No.	$Corr(\hat{r}_t, \hat{x}_{t+1})$	$Corr(\hat{x}_{t}, \hat{r}_{t+1})$	$MI(\hat{r}_t, \hat{x}_{t+1})$	$MI(\hat{x}_t, \hat{r}_{t+1})$	$TE\left(\hat{r}\rightarrow\hat{x} ight)$	$TE\left(\hat{x}\rightarrow\hat{r} ight)$
	of			(*********************************	(,	$\tau = 25$	$\tau = 25$
	Stocks						
MERVAL	10	$0.036 \pm$	$0.029 \pm$	$0.015 \pm$	$0.005 \pm$	$0.003 \pm$	$0.001 \pm$
		0.013	0.007	0.003	0.002	0.002	0.001
ASX	176	-0.008 ±	$0.005 \pm$	$0.027 \pm$	$0.026 \pm$	$0.010 \pm$	$0.009 \pm$
		0.004	0.002	0.003	0.003	0.001	0.001
BOVESPA	50	$-0.004 \pm$	0.014 ±	$0.009 \pm$	$0.008 \pm$	$0.007 \pm$	$0.004 \pm$
		0.005	0.003	0.001	0.002	0.001	0.001
TSX	216	-0.009 ±	$0.007 \pm$	$0.026 \pm$	0.024 ±	$0.009 \pm$	$0.008 \pm$
		0.002	0.001	0.002	0.002	0.001	0.001
SHANGHA	I 846	0.206 ±	$0.003 \pm$	$0.047 \pm$	0.017 ±	$0.000 \pm$	$0.005 \pm$
		0.001	0.001	0.001	0.000	0.000	0.000
CAC	40	-0.038 ±	$0.002 \pm$	0.018 ±	$0.009 \pm$	$0.004 \pm$	$0.003 \pm$
		0.004	0.002	0.002	0.002	0.001	0.001
DAX	30	-0.034 ±	$-0.004 \pm$	$0.017 \pm$	0.010 ±	$0.007 \pm$	$0.004 \pm$
		0.005	0.003	0.002	0.002	0.001	0.001
AEX	22	-0.034 ±	$0.000 \pm$	$0.024 \pm$	0.015 ±	$0.002 \pm$	$0.003 \pm$
		0.007	0.003	0.004	0.003	0.001	0.001
HSI	43	$0.026 \pm$	$0.006 \pm$	$0.028 \pm$	$0.010 \pm$	$0.001 \pm$	$0.003 \pm$
		0.007	0.004	0.002	0.001	0.001	0.000
SENSEX	28	$0.052 \pm$	$0.007 \pm$	$0.019 \pm$	$0.008 \pm$	$0.000 \pm$	$0.002 \pm$
		0.008	0.004	0.002	0.001	0.001	0.001

MIB	33	$-0.004 \pm$	$-0.007 \pm$	0.013 ±	$0.006 \pm$	$0.005 \pm$	$0.003 \pm$
		0.007	0.003	0.002	0.002	0.001	0.001
NKY	214	$0.087 \pm$	$0.004 \pm$	$0.024 \pm$	0.016 ±	$0.008 \pm$	$0.004 \pm$
		0.004	0.001	0.001	0.001	0.000	0.000
KOSPI	620	0.113 ±	-0.018 ±	$0.028 \pm$	0.018 ±	$0.002 \pm$	$0.004 \pm$
		0.002	0.001	0.001	0.000	0.000	0.000
MEXBOL	28	$0.002 \pm$	$0.017 \pm$	$0.009 \pm$	$0.005 \pm$	$0.003 \pm$	$0.002 \pm$
		0.007	0.005	0.002	0.001	0.001	0.000
IBEX	30	-0.011 ±	-0.003 ±	0.011 ±	$0.005 \pm$	$0.001 \pm$	$0.002 \pm$
		0.005	0.003	0.002	0.001	0.001	0.001
OMX	28	-0.014 ±	$0.007 \pm$	$0.012 \pm$	$0.006 \pm$	$0.004 \pm$	$0.002 \pm$
		0.005	0.004	0.002	0.001	0.001	0.000
SMI	18	-0.037 ±	$0.009 \pm$	$0.021 \pm$	$0.010 \pm$	$0.003 \pm$	$0.003 \pm$
		0.009	0.005	0.004	0.003	0.001	0.001
FTSE	93	-0.020 ±	$0.005 \pm$	0.011 ±	$0.007 \pm$	$0.003 \pm$	$0.003 \pm$
		0.003	0.002	0.001	0.001	0.000	0.000
NDX	87	-0.017 ±	$0.002 \pm$	$0.033 \pm$	$0.021 \pm$	$0.004 \pm$	$0.008 \pm$
		0.005	0.003	0.002	0.002	0.001	0.001
SPX	454	-0.036 ±	$0.008 \pm$	$0.026 \pm$	0.018 ±	$0.005 \pm$	$0.006 \pm$
		0.002	0.001	0.001	0.001	0.000	0.000

Shows the relationship between returns and volume for stocks across a range of global indices. The *Corr* $(\hat{r}_t, \hat{x}_{t+1})$ is generally negative for most of the western market stocks but insignificant or positive for eastern and emerging market stocks. Stocks on the Shanghai index exhibit a strongly positive correlation which is an order of magnitude larger than other stocks. This is probably because of the t + 1 trading rule in China which prevents investors from buying and selling the same stock on the same day. The *Corr* $(\hat{x}_t, \hat{r}_{t+1})$ is generally insignificant except in Central and Latin America where it is stronger than the *Corr* $(\hat{r}_t, \hat{x}_{t+1})$. Generally, $MI(\hat{r}_t, \hat{x}_{t+1}) > MI(\hat{x}_t, \hat{r}_{t+1})$ but the *TE* shows that the information transfer is either bi-directional or statistically insignificant.

B.2.2 Volume and Volatility

Table B.3: Summary of the Stock Level Volume-VolatilityRelation by Index

Index	No.	$Corr(\hat{x}_t, \hat{v_{t+1}})$	$Corr(\hat{v}_t, \hat{x}_{t+1})$	$MI(\hat{x}_t, \hat{v}_{t+1})$	$MI(\hat{v_t}, \hat{x}_{t+1})$	$TE\left(\hat{x}\rightarrow\hat{v} ight)$	$TE\left(\hat{v}\rightarrow\hat{x}\right)$
	of					$\tau = 25$	$\tau = 25$
	Stocks						
MERVAL	10	0.069 ±	0.123 ±	0.006 ±	0.015 ±	0.002 ±	0.000 ±
		0.025	0.034	0.003	0.003	0.001	0.001
ASX	176	$0.030 \pm$	$0.088 \pm$	$0.025 \pm$	$0.027 \pm$	0.012 ±	$0.008 \pm$
		0.006	0.008	0.003	0.002	0.001	0.001
BOVESPA	50	0.016 ±	$0.072 \pm$	$0.008 \pm$	0.010 ±	$0.005 \pm$	$0.006 \pm$
		0.008	0.013	0.002	0.002	0.001	0.001
TSX	216	$0.034 \pm$	$0.084 \pm$	$0.022 \pm$	0.024 ±	0.010 ±	$0.007 \pm$
		0.004	0.006	0.002	0.002	0.001	0.001
SHANGHA	I 846	0.129 ±	0.210 ±	0.018 ±	$0.029 \pm$	0.006 ±	$0.000 \pm$
		0.002	0.003	0.000	0.000	0.000	0.000
CAC	40	0.111 ±	0.209 ±	$0.010 \pm$	0.018 ±	$0.005 \pm$	$0.004 \pm$
		0.013	0.013	0.002	0.002	0.001	0.001
DAX	30	$0.099 \pm$	0.208 ±	0.013 ±	0.018 ±	$0.008 \pm$	$0.005 \pm$
		0.018	0.016	0.002	0.003	0.001	0.001
AEX	22	0.103 ±	$0.209 \pm$	0.015 ±	$0.025 \pm$	$0.006 \pm$	$0.003 \pm$
		0.017	0.018	0.003	0.004	0.002	0.001
HSI	43	0.117 ±	0.210 ±	$0.008 \pm$	$0.025 \pm$	$0.002 \pm$	$0.001 \pm$
		0.011	0.013	0.001	0.002	0.001	0.000
SENSEX	28	0.122 ±	0.210 ±	$0.010 \pm$	0.019 ±	$0.002 \pm$	$0.000 \pm$
		0.012	0.016	0.002	0.002	0.001	0.001
MIB	33	$0.085 \pm$	0.209 ±	$0.007 \pm$	0.015 ±	$0.004 \pm$	$0.004 \pm$
		0.014	0.015	0.002	0.003	0.001	0.001

NKY	214	$0.091 \pm$	$0.182 \pm$	$0.014 \pm$	0.019 ±	$0.005 \pm$	$0.006 \pm$
		0.004	0.005	0.001	0.001	0.000	0.000
KOSPI	620	0.163 ±	$0.255 \pm$	$0.019 \pm$	$0.027 \pm$	$0.005 \pm$	$0.003 \pm$
		0.004	0.005	0.001	0.001	0.000	0.000
MEXBOL	28	$0.033 \pm$	$0.065 \pm$	$0.004 \pm$	$0.007 \pm$	$0.003 \pm$	$0.003 \pm$
		0.008	0.013	0.001	0.001	0.001	0.001
IBEX	30	$0.055 \pm$	$0.099 \pm$	$0.007 \pm$	0.011 ±	$0.002 \pm$	$0.001 \pm$
		0.009	0.012	0.001	0.002	0.001	0.000
OMX	28	$0.067 \pm$	0.138 ±	$0.006 \pm$	$0.012 \pm$	$0.003 \pm$	$0.003 \pm$
		0.012	0.015	0.001	0.002	0.001	0.000
SMI	18	0.128 ±	0.175 ±	0.013 ±	$0.023 \pm$	$0.004 \pm$	$0.002 \pm$
		0.023	0.024	0.002	0.004	0.001	0.001
FTSE	93	$0.060 \pm$	0.116 ±	$0.007 \pm$	$0.010 \pm$	$0.003 \pm$	$0.002 \pm$
		0.006	0.007	0.001	0.001	0.000	0.000
NDX	87	0.102 ±	0.187 ±	$0.022 \pm$	$0.034 \pm$	$0.009 \pm$	$0.004 \pm$
		0.007	0.010	0.002	0.002	0.001	0.001
SPX	454	$0.092 \pm$	0.177 ±	$0.015 \pm$	$0.023 \pm$	$0.006 \pm$	$0.004 \pm$
		0.003	0.004	0.001	0.001	0.000	0.000

Shows the relationship between volume and volatility for stocks across a range of global indices. The $Corr(\hat{x}_t, \hat{v}_{t+1})$ and $Corr(\hat{v}_t, \hat{x}_{t+1})$ are both strongly positive with the latter being stronger; this is consistent with the *MI* results. The *TE* shows that the information transfer is either bi-directional or statistically insignificant.

B.2.3 Returns and Volatility

Table B.4: Summary of the Stock Level Return-VolatilityRelation by Index

Index	No.	$Corr(\hat{r}_t, \hat{v}_{t+1})$	$Corr(\hat{v}_t, \hat{r}_{t+1})$	$MI(\hat{r}_t, \hat{v}_{t+1})$	$MI(\hat{v}_{t},\hat{r}_{t+1})$	$TE\left(\hat{r}\rightarrow\hat{v}\right)$	$TE\left(\hat{v}\rightarrow\hat{r} ight)$
	of					$\tau = 25$	au = 25
	Stocks						
MERVAL	10	-0.045 ±	0.074 ±	0.016 ±	0.017 ±	0.001 ±	0.004 ±
		0.031	0.034	0.004	0.005	0.001	0.001
ASX	176	-0.031 ±	$0.007 \pm$	$0.037 \pm$	0.037 ±	$0.005 \pm$	$0.004 \pm$
		0.004	0.004	0.004	0.004	0.001	0.000
BOVESPA	50	-0.048 ±	$0.049 \pm$	0.011 ±	0.011 ±	0.001 ±	$0.002 \pm$
		0.008	0.006	0.001	0.001	0.000	0.000
TSX	216	-0.021 ±	-0.002 ±	$0.030 \pm$	$0.029 \pm$	$0.003 \pm$	$0.003 \pm$
		0.004	0.004	0.002	0.002	0.000	0.000
SHANGHA	I 846	$-0.005 \pm$	0.014 ±	$0.007 \pm$	$0.008 \pm$	$0.000 \pm$	$0.002 \pm$
		0.002	0.002	0.000	0.000	0.000	0.000
CAC	40	-0.043 ±	0.016 ±	$0.007 \pm$	$0.006 \pm$	$0.000 \pm$	$0.002 \pm$
		0.006	0.005	0.001	0.001	0.000	0.000
DAX	30	$-0.050 \pm$	$0.003 \pm$	$0.008 \pm$	$0.008 \pm$	$0.001 \pm$	$0.004 \pm$
		0.006	0.006	0.001	0.001	0.000	0.000
AEX	22	-0.043 ±	0.001 ±	$0.012 \pm$	$0.012 \pm$	$0.001 \pm$	$0.003 \pm$
		0.009	0.007	0.002	0.002	0.001	0.001
HSI	43	-0.064 ±	$0.025 \pm$	0.013 ±	0.013 ±	$0.003 \pm$	$0.004 \pm$
		0.009	0.007	0.001	0.001	0.001	0.000
SENSEX	28	-0.022 ±	0.011 ±	$0.008 \pm$	$0.008 \pm$	$0.000 \pm$	$0.001 \pm$
		0.007	0.008	0.001	0.001	0.000	0.000
MIB	33	-0.038 ±	0.015 ±	0.011 ±	$0.009 \pm$	$0.000 \pm$	$0.003 \pm$
		0.008	0.006	0.004	0.004	0.001	0.000

NKY	214	-0.031 ±	$0.032 \pm$	$0.015 \pm$	$0.015 \pm$	$0.002 \pm$	$0.002 \pm$
		0.002	0.002	0.001	0.001	0.000	0.000
KOSPI	620	0.013 ±	$0.007 \pm$	$0.032 \pm$	$0.032 \pm$	$0.002 \pm$	$0.003 \pm$
		0.002	0.002	0.001	0.001	0.000	0.000
MEXBOL	28	-0.038 ±	0.011 ±	0.015 ±	$0.014 \pm$	$0.000 \pm$	$0.002 \pm$
		0.013	0.013	0.002	0.002	0.001	0.000
IBEX	30	$-0.023 \pm$	$0.008 \pm$	$0.008 \pm$	$0.008 \pm$	$0.000 \pm$	$0.002 \pm$
		0.008	0.006	0.001	0.001	0.001	0.001
OMX	28	$-0.027 \pm$	0.018 ±	$0.010 \pm$	$0.009 \pm$	$0.001 \pm$	$0.003 \pm$
		0.007	0.009	0.001	0.001	0.001	0.001
SMI	18	-0.041 ±	$0.008 \pm$	$0.010 \pm$	$0.010 \pm$	$0.001 \pm$	$0.003 \pm$
		0.012	0.013	0.002	0.002	0.001	0.001
FTSE	93	-0.031 ±	0.011 ±	0.011 ±	$0.010 \pm$	$0.001 \pm$	$0.003 \pm$
		0.005	0.005	0.001	0.001	0.000	0.000
NDX	87	-0.030 ±	0.018 ±	0.016 ±	0.016 ±	$0.003 \pm$	$0.004 \pm$
		0.005	0.005	0.001	0.001	0.000	0.000
SPX	454	-0.047 ±	$0.022 \pm$	0.014 ±	0.013 ±	$0.002 \pm$	$0.003 \pm$
		0.002	0.002	0.001	0.001	0.000	0.000

Shows the relationship between return and volatility for stocks across a range of global indices. The Leverage Effect is clearly identifiable in all the markets by the negative $Corr(\hat{r}_t, \hat{v}_{t+1})$, except for the South Korean (KOSPI) stocks. However, the $Corr(\hat{v}_t, \hat{r}_{t+1})$ is positive. $MI(\hat{r}_t, \hat{v}_{t+1}) \approx MI(\hat{v}_t, \hat{r}_{t+1})$ but the TE is not statistically significant over this period.

Controlling for Trading Volumes and Correlation

Table B.5: Effective Mutual Information for the Stock Level Leverage Effect by Index Controlling for Trading Volumes and Index Returns

Index	No.	$EMI(\hat{r}_t, \hat{v}_{t+1})$	$EMI(\hat{r}_t, \hat{v}_{t+1} \mid \hat{x}_t, \hat{x}_{t+1})$	$EMI(\hat{r}_t, \hat{v}_{t+1} \hat{x}_t, \hat{x}_{t+1}, \hat{r}_{I,t}, \hat{r}_{i,t+1})$
	of			
	Stocks			
MERVAL	10	0.012 ± 0.004	0.005 ± 0.001 (42%)	0.000 ± 0.001 (0%)
ASX	176	0.033 ± 0.004	$0.018 \pm 0.002 \ (55\%)$	0.006 ± 0.001 (18%)
BOVESPA	50	0.007 ± 0.001	0.006 ± 0.001 (86%)	$0.001 \pm 0.000 (14\%)$
TSX	216	0.026 ± 0.002	$0.012 \pm 0.001 \ (46\%)$	$0.004 \pm 0.000 (15\%)$
SHANGHA	I 846	0.003 ± 0.000	0.003 ± 0.000 (100%)	$0.000 \pm 0.000 \ (0\%)$
CAC	40	0.003 ± 0.001	$0.005 \pm 0.001 \ (167\%)$	0.002 ± 0.000 (67%)
DAX	30	0.004 ± 0.001	0.006 ± 0.001 (150%)	$0.001 \pm 0.001 \ (25\%)$
AEX	22	0.008 ± 0.002	0.005 ± 0.001 (63%)	$0.000 \pm 0.001 \ (0\%)$
HSI	43	0.009 ± 0.001	$0.004 \pm 0.001 (44\%)$	$0.000 \pm 0.000 \ (0\%)$
SENSEX	28	0.004 ± 0.001	$0.001 \pm 0.001 \ (25\%)$	0.001 ± 0.001 (25%)
MIB	33	0.007 ± 0.004	0.008 ± 0.002 (114%)	$0.001 \pm 0.001 \ (14\%)$
NKY	214	0.011 ± 0.001	0.007 ± 0.000 (64%)	0.004 ± 0.000 (36%)
KOSPI	620	0.028 ± 0.001	0.019 ± 0.001 (68%)	$0.007 \pm 0.000 \ (25\%)$
MEXBOL	28	0.011 ± 0.002	0.008 ± 0.002 (73%)	0.002 ± 0.001 (18%)
IBEX	30	0.004 ± 0.001	$0.004 \pm 0.001 \ (100\%)$	$0.000 \pm 0.001 \ (0\%)$
OMX	28	0.006 ± 0.001	0.005 ± 0.001 (83%)	$0.001 \pm 0.000 (17\%)$
SMI	18	0.006 ± 0.002	0.004 ± 0.001 (67%)	$0.000 \pm 0.001 \ (0\%)$
FTSE	93	0.007 ± 0.001	0.006 ± 0.000 (86%)	$0.001 \pm 0.000 (14\%)$
NDX	87	0.012 ± 0.001	$0.006 \pm 0.001 \ (50\%)$	$0.002 \pm 0.000 (17\%)$
SPX	454	0.009 ± 0.001	$0.005 \pm 0.000 \ (56\%)$	$0.001 \pm 0.000 \ (11\%)$

Shows the $EMI(\hat{r}_t, \hat{v}_{t+1})$ and the $EMI(\hat{r}_t, \hat{v}_{t+1}|Z)$ controlling for trading volumes, \hat{x} , and index returns, \hat{r}_i , for stocks across a range of global indices. The percentage of unexplained EMI is shown in brackets. It shows that most of the EMI can be explained by trading volumes and index returns. The impact of trading volumes varies across markets but on average it appears to account for 40-50% of the EMI.

Appendix C

An Information Theoretic Analysis of Returns, Volatility and Trading Volumes for Emerging Market Indices

In Chapter 5 I examined the dynamic relationships between returns, volatility and trading volumes for developed market indices. Here I extend this analysis to a range of emerging market indices.

C.0.4 The Data

The data is sourced from Bloomberg and covers daily returns and volumes for developed stock indices covering the period 1980-2012. The indices include Argentina (MERVAL), Brazil (BOVESPA), China (SHANGHAI), India (SENSEX) and Mexico (MEXBOL). The data has not been corrected for stocks that have been added/removed from the indices. However, I do not believe that this unduly affects the results.

All statistics and results have been calculated at the individual index level and then averaged to give a mean value for all indices. The significance levels have been estimated by calculating the various measures using surrogate data sets with similar statistical properties but without the inter-relationships; this is consistent with similar research.

C.1 Emerging Markets Results

C.1.1 Persistence of Returns, Volatility and Volumes



Figure C.1: Auto-Mutual Information and Auto-Covariance Functions for Emerging Market Indices at the Daily Frequency

Shows the auto-mutual information (Top) and normalised auto-covariance (Bottom) functions. They are calculated for returns (blue stars with dashed lines), volatility (green squares with dot-dashed lines) and volumes (red diamonds with solid lines). These are given with associated one standard errors, exponential curve fits and 95% Significance Levels (black dashed lines). The *MI* is statistically significant and persistent for volumes but it is only statistically significant for 4 days for volatility and not at all for returns. The auto-covariance function shows that both volumes and volatility are statistically significant and persistent but returns are not statistically significant at any time horizon.

C.1.2 Returns and Volume



Figure C.2: Cross-Mutual Information and Cross-Covariance Functions for Emerging Market Index Returns and Volume at the Daily Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions between returns and volume. The blue squares with solid lines show the $MI(\hat{x}_t, \hat{r}_{t+\tau})$ and $Corr(x_t, r_{t+\tau})$. The green diamonds with dot-dashed lines show the $MI(\hat{r}_t, \hat{x}_{t+\tau})$ and $Corr(r_t, x_{t+\tau})$. These are given with associated one standard errors, exponential curve fits and 95% Significance Levels (dashed black lines). The *MI* is statistically significant and persistent which implies bi-directional causality. This contrasts with the cross-covariance function which only shows $Corr(r_t, x_{t+1})$ as statistically significant. Interestingly, the correlation is positive which contrasts with previous results. This is likely because of the t + 1 trading rule on the Shanghai stock market which prevents investors from buying and selling the same stock on the same day.



Figure C.3: Cross-Partial Mutual Information and Cross-Transfer Entropy Functions for Emerging Market Index Returns and Volume at the Daily Frequency

(Top) Shows the partial cross-mutual information function for returns and volume, where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{x}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised returns. It is given with associated one standard errors, exponential curve fits and a 95% Significance Level (dashed black line). The *PMI* is statistically significant in both directions for $\tau < 3$ which bi-directional causality. Since neither *MI* is statistically significant for $\tau > 5$, the *PMI* also implies that the persistence is due to auto-information. (Bottom) Shows the cross-transfer entropy function for returns and volumes. The green diamonds represent $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ and the blue squares represent the $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ and the dashed blue line for $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau})$. The *TE* shows that $TE(\hat{x}_t \rightarrow \hat{r}_{t+\tau}) > TE(\hat{r}_t \rightarrow \hat{x}_{t+\tau})$ for $\tau < 5$ and statistically significant but equal for $\tau > 5$. Again this implies bi-directional (Granger) causality.

C.1.3 Volume and Volatility



Figure C.4: Cross-Mutual Information and Cross-Covariance Functions for Emerging Market Index Volatility and Volume at the Daily Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for volumes and volatility. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau})$ and the $Cov(v_t, x_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau})$ and the $Cov(x_t, v_{t+\tau})$. Both graphs are given with associated one standard errors, exponential curve fits and 95% Significance Levels (black dashed lines). The MI is statistically significant in both directions with $MI(\hat{v}_t, \hat{x}_{t+1}) \approx MI(\hat{x}_t, \hat{v}_{t+1})$. This indicates bi-directional causality between volume and volatility. However, the cross-covariance is only statistically significant for 1 day in either direction.


Figure C.5: Cross-Partial Mutual Information and Cross-Transfer Entropy Functions for Emerging Market Index Volatility and Volume at the Daily Frequency

(Top) Shows the partial cross-mutual information function for volumes and volatility where I have controlled for auto-information. The green diamonds with dot-dashed lines represent the $MI(\hat{v}_t, \hat{x}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volumes. The blue squares with solid lines represent the $MI(\hat{x}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volatility. It is given with associated one standard errors, exponential curve fits and a 95% Significance Level (black dashed line). The MI is only statistically significant for 1-2 days in both directions This indicates bi-directional causality and since neither MI is statistically significant for $\tau > 2$, the PMI also implies that the persistence is due to auto-information. (Bottom) Shows the cross-transfer entropy function for volumes and volatility. The green diamonds represent $TE(\hat{v}_t \rightarrow \hat{x}_{t+\tau})$ and the blue squares represent $TE(\hat{x}_t \rightarrow \hat{v}_{t+\tau})$. This is given with associated one standard errors and 95% Significance Levels where the green dot-dashed line is for the $TE(\hat{v}_t \rightarrow \hat{x}_{t+\tau})$ and the dashed blue line is for the $TE(\hat{x}_t \rightarrow \hat{v}_{t+\tau})$. The results for the TE are mixed but overall it appears that there is bi-directional (Granger) causality between volume and volatility.

C.1.4 Returns and Volatility



Figure C.6: Cross-Mutual Information and Cross-Covariance Functions for Emerging Market Index Returns and Volatility at the Daily Frequency

Shows the cross-mutual information (Top) and normalised cross-covariance (Bottom) functions for returns and volatility. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau})$ and the $Cov(r_t, v_{t+\tau})$. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau})$ and the $Cov(v_t, r_{t+\tau})$. Both graphs are given with associated one standard errors, exponential curve fits and 95% Significance Levels (black dashed lines). The *MI* only appears statistically significant in either direction for a few days and of equal magnitude (within the margin of error). This implies a bi-directional information flow between returns and volatility. However, the relationship is much clearer in the cross-covariance function where the Leverage Effect is statistically significant for around 5 days.



Figure C.7: Cross-Partial Mutual Information and Cross-Transfer Entropy Functions for Emerging Market Index Returns and Volatility at the Daily Frequency

(Top) Shows the partial cross-mutual information function for returns and volatility. The green diamonds with dot-dashed lines represent the $MI(\hat{r}_t, \hat{v}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised volatilities. The blue squares with solid lines represent the $MI(\hat{v}_t, \hat{r}_{t+\tau}|Z)$ where $Z = [\hat{j}_t, ..., \hat{j}_{t+\tau-1}]$ for $\tau \ge 1$ and \hat{j} are the normalised returns. These are given with associated one standard errors, exponential curve fits and a 95% Significance Level (dashed black line). As with the cross-mutual information function there is evidence of bi-directional information flow but the results are not clear. (Bottom) Shows the cross-transfer entropy functions for returns and volatility. The green diamonds represent the $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the blue squares represent the $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. These are given with associated one standard errors and 95% Significance Levels which are represented by the green dot-dashed line for $TE(\hat{r}_t \rightarrow \hat{v}_{t+\tau})$ and the dashed blue line for $TE(\hat{v}_t \rightarrow \hat{r}_{t+\tau})$. The *TE* gives evidence of bi-directional (Granger) causality between returns and volatility but it is not clear.

Controlling for Trading Volumes and Correlation

Table C.1: Mutual Information for the Emerging Market Index Level Leverage Effect Controlling for Trading Volumes and Correlation Asymmetry

	$MI\left(\hat{r},\hat{v} ight)$	$MI(\hat{r},\hat{v}\mid\hat{x})$	$MI(\hat{r},\hat{v}\mid\hat{c},)$	$MI(\hat{r},\hat{v}\mid\hat{x},\hat{c},)$	
(r_t, v_{t+1})	0.004 ± 0.003	0.010 ± 0.001	0.189 ± 0.075	0.067 ± 0.017	
	(0.001)	(<0.001)	(<0.001)	(<0.001)	
(v_t, r_{t+1})	0.000 ± 0.001	0.006 ± 0.001	0.190 ± 0.074	0.067 ± 0.017	
	(1.000)	(<0.001)	(<0.001)	(<0.001)	
(r_t, v_{t+2})	0.007 ± 0.003	0.010 ± 0.001	0.174 ± 0.060	0.084 ± 0.030	
	(<0.001)	(<0.001)	(<0.001)	(<0.001)	
(v_t, r_{t+2})	0.004 ± 0.002	0.009 ± 0.002	0.180 ± 0.068	0.082 ± 0.029	
	(0.001)	(<0.001)	(<0.001)	(<0.001)	

Shows the *MI* between returns and volatility at the 1 and 2 day time lags for emerging market indices. The p-values are shown in brackets. Controlling for trading volumes alone has no effect on the *MI* between returns and volatility but correlation asymmetry is shown to dampen the *MI* by two orders of magnitude. Controlling for correlation asymmetry and trading volumes increases the *MI* between returns and volatility by an order of magnitude. These results indicate that correlation asymmetry actually dampens the *MI* between returns and volatility but once this has been accounted for trading volumes become an important driver of the Leverage Effect. This is consistent with the results for developed market indices.

Appendix D

Fitting the Multivariate q-Gaussian Distribution

In Chapter 8 I showed how multivariate q-Gaussian distributions could be used to model the IVS for Goldman Sachs. Figure D.1 shows the fitted marginal distributions for returns and 3, 6, 12 and 18 month implied volatility for several large US stocks. It is clear that the q-Gaussian distribution (blue lines) fits the data far better than the Gaussian distribution (black lines). Qualitatively it also appears to fit the tails of the distribution very well. These results are consistent with those for Goldman Sachs.



Figure D.1: Univariate q-Gaussian Fits of Returns and Implied Volatility Changes for US Stocks

This shows the fitted marginal distributions for returns and 3, 6, 12 and 18 month implied volatility for General Electric, Microsoft, Proctor and Gamble and Exxon Mobile (Left to Right). It is clear that the q-Gaussian distribution fits (blue lines) fit the data far better than the Gaussian distribution. Qualitatively it also appears to fit the tails of the distribution very well. These results are consistent with those shown previously for Goldman Sachs.

Calculating the AIC for the q-Gaussian model directly requires knowing the normalising constant, which is difficult to calculate in high dimensions. Instead, we use the fact that a q-Gaussian with $\beta = 0.5$, $q \rightarrow 1$ becomes the Gaussian distribution; the loglikelihood difference between a q-Gaussian and a Gaussian can then be calculated with Equation 8.7, where $\phi = (0.5, 1 + \varepsilon)$ and $\varepsilon \ll 1$. This log-likelihood difference can then be inserted into Equation 8.8.

We can test this method for calculating the AIC difference by applying it to various data sets simulated from q-Gaussian models with different values of q, shown in Figure D.2. For values of q close to 1 the AIC favours the Gaussian distribution - it has fewer parameters, and the q-Gaussian fit is very close to a Gaussian. As q increases and the simulated data becomes heavier tailed, the q-Gaussian fit becomes better and the AIC difference drops below 0, favouring it over the Gaussian.



Figure D.2: Akaike Information Criterion difference between the Gaussian and q-Gaussian Models

This shows the AIC difference between the Gaussian and q-Gaussian models when fitted to simulated data with different values of q. For values of q close to 1 the AIC favours the Gaussian distribution - it has fewer parameters, and the q-Gaussian fit is very close to a Gaussian. As q increases and the simulated data becomes heavier tailed, the q-Gaussian fit becomes better and the AIC difference drops below 0, favouring it over the Gaussian.