University of Alberta

Essays on Extension of Trading Time and Value at Risk

by

Ebenezer Asem ©

A thesis submitted to the **Faculty of Graduate Studies and Research** in partial fulfillment of the requirement for the degree of **Doctor of Philosophy**

in

FINANCE

Faculty of Business

Edmonton, Alberta

Spring 2002

0-612-68538-1

University of Alberta

Library Release Form

Name of Author: Ebenezer Asem

Title of Thesis: Essays on Extension of Trading Time and Value at Risk

Degree: Doctor of Philosophy

Year this Degree Granted: 2002

Permission is hereby granted to the University of Alberta Library to reproduce single copies of this thesis and to lend or sell such copies for private, scholarly, or scientific research purposes only.

The author reserves all other publication and other rights in association with the copyright in the thesis, and except as hereinbefore provided, neither the thesis nor any substantial portion thereof may be printed or otherwise reproduced in any material form whatever without the author's prior written permission.

EAsem

Ebenezer Asem
Department of Finance and Management Science
Faculty of Business
University of Alberta
Edmonton, T6G 2R6

Date: Dec. 19, 2001

UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recommend to the Faculty of Graduate Studies and Research for acceptance, a thesis entitled Essays on Extension of Trading Time and Value at Risk submitted by Ebenezer Asem in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Finance.

Dr. S. Beveridge

Dr. A. Kaul

Dr. V. Mehrotra

Dr. J. Unterschultz

Dr. H. H. Zhang

October 26, 2001

DEDICATION

To my parents

ABSTRACTS

I. The NYSE extended its trading hours on October 1, 1974, and also on September 30, 1985. These events provide ideal opportunities to examine the sources of volume and return variability, an issue that has become especially relevant given the current move towards continuous trading. We find that the extension of trading time in 1974 did not increase trading activity but trading activity increases after the extension in 1985. These results are consistent with our proposed information cancellation hypothesis. Our finding has useful implication for extending trading periods. It suggests that extending trading time to periods when businesses are closed and information arrival is low may not generate significant increases in trading activity. The study also shows that extending trading hours would not increase return variability. In addition, we find that extending trading time reduces transitory noise in opening prices relative to closing prices, and the extension changes intraday return variances which reflects changes in the arrival of private traders. The latter finding is consistent with the price formation hypothesis and the former supports the private information hypothesis.

II. An accurate estimation of Value at Risk (VaR) requires proper modeling of the unconditional kurtosis of the risk factors as well as appropriate apportioning of the modeled kurtosis between stochastic volatility and the distribution of the risk factors. In GARCH models, the division of the unconditional kurtosis between time varying variances and the distribution is determined by the assumed conditional distribution of the errors. We examine the importance of this by applying normal and Student's t-distributions' filtered historical simulations to five major exchange rates. The study shows

that the accuracy of VaR estimates of the British Pound can be improved by using appropriate fat-tailed distributions rather than more general stochastic volatility models. This finding suggests that, for some risk factors, the source of the empirical kurtosis is crucial in appropriately modeling the future distribution of the risk factors. In addition, we find that, for the purposes of forecasting VaRs of direct exposures, a more pertinent measure of kurtosis is the number of standard deviations associated with the particular confidence level.

ACKNOWLEDGEMENT

I would like to express my sincere thanks to my dissertation committee members, Dr. Steve Beveridge, Dr. Aditya Kaul, and Dr. Vikas Mehrotra, to the external committee member, Dr. Jim Unterschultz, and to the external examiner, Dr. Harold Zhang for the time and effort they have expended on my behalf.

I am particularly indebted to my supervisors, Dr. Steve Beveridge and Dr. Aditya Kaul, who did not only direct and supervise this thesis but also have profound impact on my life throughout my studies. I will also like to thank Dr. G. Barone-Adesi for his help and useful suggestions at the initial stages of the thesis. Finally, I owe special thanks to the Ph.D. finance students, especially Gloria Tian, for always being there.

I am very grateful for the financial assistance provided by Mutual Life and Domtar Ph.D. Fellowship to support my study.

I cannot end my acknowledgement properly without special thanks to my wife, Edna Djokoto, and my daughter, Eyram Asem, for their support during the program.

TABLE OF CONTENT

	Page
CHAPTER 1: INTRODUCTION	1
CHAPTER 2: THE SOURCES OF VOLUME AND RETURN VARIABILITY: EVIDENCE FROM EXTENSION OF NYSE TRADING HOURS	
2.1. INTRODUCTION	3
2.2. Review of Theories and Testable Hypotheses	7
2.2.1. The Effects of the Early Opening on Daily Volume and Return Variability	8
2.2.2. The Effects of the Early Opening on some Empirical Facts	12
2.2.2.1. Intraday Volume and Return exhibit U-shaped pattern	13
2.2.2.2. Hourly Returns are more Volatile during Trading Time than Non-Trading Time	15
2.2.2.3. Open-to-Open Returns have more Transitory Volatility Close-to-Close Returns	17
2.3. DATA DESCRIPTION	20
2.4. EMPIRICAL RESULTS: THE EARLY OPENING AND TRADE	
2.4.1. Regression Analysis	22
2.4.2. Accounting for Volume in the New Period	27
2.4.3. Other Evidence	32
2.5. EXPLANATIONS FOR THE INCREASE IN TRADE POST-EARLY OPENING	34
2.5.1. Private Information Hypothesis	36
2.5.2. Noise Trading Hypothesis	39
2.5.3. Lower Overnight Risk Hypothesis	40
2.5.4. Information Cancellation Hypothesis	42
2.6. EFFECTS OF THE EARLY OPENING ON SOME EMPIRICAL OBSERVATIONS	
2.6.1. Intraday Volume and Return Variability	44
2.6.2. Trading versus Non-Trading Return Variability	45
2.6.3. Opening versus Closing Transitory Noise	47
2.7. CONCLUSIONS	49
Tables	52
Figures	58

TABLE OF CONTENT- CONTINUED

References 61

Appendix. A Model of Information Flow and Volume of Trade 63

CHAPTER 2: KURTOSIS AND THE ACCURACY OF VALUE AT RISK

3.1. INTRODUCTION 67

3.2. MOTIVATION AND BACKGROUND INFORMATION
- 3.2.1. Review of Value at Risk 72
- 3.2.2. Some Empirical Measurement Issues 74

3.3. EVALUATION PROCEEDURES
- 3.3.1. Operational Evaluation 79
- 3.3.2. Observed and Predicted Exceedences 79

3.4. DATA AND EMPIRICAL ESTIMATIONS 84
- 3.4.1. Summary Statistics of of Exchange Rate Returns 85

3.5. THE GARCH MODELS 88
- 3.5.1. The Normal Likelihood
  - 3.5.1.1. Estimation 88
  - 3.5.1.2. Forecasting 91
- 3.5.2. The t-distribution Likelihood
  - 3.5.2.1. Estimation 101
  - 3.5.2.2. Forecasting 106

3.6. DISCUSSION
- 3.6.1. The Theoretical Forecasts 117
- 3.6.2. The Empirical Forecasts 121
- 3.6.3. The Theoretical versus the Empirical Forecasts 124

3.7. CONCLUSIONS 125

Tables 128

Figures 140

References 141

CHAPTER 4: CONCLUSION 144

LIST OF TABLES

Table 2.1: Regression of Daily Turnover (with and without trend square): Two Years Around the Early Opening 52

Table 2.2: Difference in Mean Turnover Tests: Three Month Before and After September 30th for the Years 1980 to 1990 53

Table 2.3: Regressions of Daily Turnover of Stocks listed on the LSE and the NYSE and of Stocks listed on the NYSE but *not* the LSE 54

Table 2.4: Changes in the One-And-Half Hour turnover after the Early-Opening 55

Table 2.5: Variances and Relative Variances of One-and-Half Hour S&P 500 Returns 55

Table 2.6: Variance Ratios to Estimate the Correlations among Intraday Returns 56

Table 2.7: Variances and Relative Variances of Thirty-Minute S&P 500 Returns Two Years before the Early Opening and Two Years after it. 56

Table 2.8: Regression of Daily Turnover of all Stocks: Two Years Around the Late Closure 57

Table 3.1: Summary Statistics of Daily Exchange Rate Returns (In-Sample) 128

Table 3.2: Summary Statistics of Daily Exchange Rate Returns (Out-of-Sample) 128

Table 3.3: Autocorrelations and Daily Exchange Rate Returns and Squared Returns (In-Sample) 129

Table 3.4: GARCH Estimates (Normal Likelihood) 129

Table 3.5: Summary Statistics of Standardized Residuals (GARCH/Normal) 130

Table 3.6: Various Statistics of the Forecasts of the Normal Distribution Likelihood (Empirical and Theoretical Forecasts) 131

Table 3.7: Likelihood Ratio Test Statistics of the Normal Forecasts 132

Table 3.8: Selected and Estimated t-distributions GARCH Estimates 133

LIST OF TABLE - CONTINUED

Table 3.9: Summary Statistics of the Standardized Residuals from the t-distributions estimates (In-Sample) 134

Table 3.10: Various Statistics of the Forecasts of the Selected t-distribution Likelihood (Empirical and Theoretical Forecasts) 135

Table 3.11: Various Statistics of the Forecasts of the Estimated t-distribution Likelihood (Empirical and Theoretical Forecasts) 136

Table 3.12: Likelihood Ratio Test Statistics of the Selected t-distribution Forecasts 137

Table 3.13: Likelihood Ratio Test Statistics of the Estimated t-distribution Forecasts 138

Table 3.14: Standard Deviation of the Forecast Variances under the Different Likelihoods 139

Table 3.15: Average Confidence Intervals of Forecast Empirical Distributions 139

LIST OF FIGURES

Fig. 2.1: Thirty-Minute Turnovers Two Years Before the Early Opening and Two Years After it 58

Fig. 2.2: Thirty-Minute Turnover: Comparison of Trading Time Effects 59

Fig. 2.3: Overnight and Thirty-Minute Trade Time S&P 500 Return Variances 60

Fig. 3.1: Empirical Distributions of the Standardized Errors of the GARCH/t-distributions 140

CHAPTER 1

INTRODUCTION

The thesis presents two papers representing research on the effects of extending trading time and the accuracy of forecasting value-at-risk. The unifying theme of the thesis is the study of return variability. The first paper we studies the influence of extending trading time on return variability and the second studies the importance of modeling time-varying return variability in forecasting the value-at-risk.

In the second chapter, *The Sources of Volume and Return Variability: Evidence from Extension of NYSE Trading Hours*, we investigate the effect of extending trading time on trading activity and return variability. This chapter is motivated by the current move towards extending trading hours and the effects this may have on trading activity and return variability. We use evidence from the extensions of trading hours on the NYSE to study the effect of continuous trading. We find that the extension of NYSE trading time on September 30, 1985 increases trading activity, but trading activity did not increase after the extension on October 1, 1974. These results are consistent with our proposed information cancellation hypothesis, which suggests that an extension of trading time would generate significant trading activity if it substantially precludes information offsetting that occurs when the markets are closed. Our finding has useful implications for extending trading hours. It suggests that extending trading time to include periods when businesses are closed and information arrival is low may not generate significant increases in trading activity. Our study also shows that extending trading hours would not increase market return variability and, therefore, continuous trading may not influence market risk premiums. In addition, we find that extending trading time reduces transitory noise in opening prices relative to closing prices, and the extension changes intraday return variances which reflects changes in the arrival of informed traders. The former finding is consistent with the price formation hypothesis and the latter supports the private information hypothesis proposed to explain the higher trading time return variance relative to non-trading time return variance.

In the second chapter, *Kurtosis and the Accuracy of Value-at-Risk*, we study the effects of role of stochastic volatility in accurately modeling the Value at Risk. This study shows that an accurate estimation of Value at Risk (VaR) requires proper modeling of the unconditional kurtosis of the risk factors as well as appropriate apportioning of the modeled kurtosis between stochastic volatility and the distribution of the risk factors. In GARCH models, the division of the unconditional kurtosis between time varying variances and tail-fatness is determined by the assumed conditional distribution of the errors. We examine the importance of this by applying normal and Student's t-distributions' filtered historical simulations to five major exchange rates. The study shows that the accuracy of VaR estimates of the British Pound can be improved by using appropriate fat-tailed distributions rather than more general stochastic volatility models. This finding is important because it suggests that, for some risk factors, the source of the empirical kurtosis is crucial in appropriately modeling the future distribution of the risk factors. Furthermore, the study shows that, for the purposes of forecasting VaRs of direct exposures, a more pertinent measure of kurtosis is the number of standard deviations associated with the particular confidence level.

CHAPTER 2

THE SOURCES OF VOLUME AND RETURN VARIABILITY: EVIDENCE FROM THE EXTENSION OF NYSE TRADING HOURS

2.1. INTRODUCTION

On September 30, 1985, the New York Stock Exchange (NYSE) extended its trading hours by shifting its opening time from 10:00 a.m. to 9:30 a.m. (we refer to this as the *early opening*). This study exploits the changes in trading hours to examine the effects of longer trading time on market volume and return volatility. In addition, the extensions of trading time enables us to carry out tests, which shed light on several hypotheses proposed in the market microstructure literature, and assess some of the predictions of the model by Hong and Wang (2000).

Direct research on the effects of extending trading hours is sparse. The exceptions are Barclay, Litzenberger, and Warner (1990) and Booth and Chowdhury (1996).[1] These studies focus on using information from the effects of extending trading time on return variability to help explain the findings that returns are more volatile during trading periods than non-trading periods [see, for example, Fama (1965), Granger and Morgenstern (1970), Oldfield and Rogalski (1980), and French and Roll (1986)].[2] In addition to using direct evidence from the extended trading hours on return variability, Barclay et al. study the relationship between volume and return variability to further assist in identifying the reason(s) for the observed phenomenon. In this chapter, we study the effect of extending trading time on volume and, unlike Barclay et al., we focus on

[1] Booth and Chowhury study the effect of extending trading hours on return volatility, and they find evidence consistent with the public information hypothesis.
[2] Three hypotheses have been proposed for this phenomenon: private information based trading, arrival of public information during the day, and noise trading.

determining the sources of the volume in the new trading period with the view of explaining the microstructural relationship between trading time and volume. This exercise will help predict the effect of extending trading hours on volume which has become important in light of the current move towards extending trading time.[3] Apart from the different focus of our research, the new trading period on the Tokyo Stock Exchange, studied by Barclay et al., was concentrated on some Saturdays (3 hours of trade), but the increase in trading time on the NYSE is equally spread among the trading days (30 minutes each day).[4] This enables us to study the impact of extending trading hours on intraday trading activity, which is very relevant in the context of extending daily trading hours, but cannot be studied in the Barclay et al. experimental environment.

The volume associated with the new trading period could arise because of shifts of trades from other periods, shifts from other exchanges, or because the new trading period generates trade. One reason for the latter effect is that, after the extension of trading hours, investors can trade on accumulated overnight information before any offsetting information arrives between 9:30 a.m. - 10:00 a.m. Thus, to the extent that some of the information arriving while the market is closed is offset by other information (e.g., good news offset by subsequent bad news), trading volume will rise if the markets are open for longer periods.

We investigate the effect of extending trading hours on volume by using difference in mean tests and regression techniques. In addition, we study the sources of the volume in the new trading period to ascertain whether part of it is attributable to the

---

[3] Presently, both the NYSE and the NASDAQ/AMEX have established Extended Hours Working Groups to work out the modalities of extending trading hours.
[4] The three-hour increase in trading time is comparable to the two and a half hour weekly gain on the NYSE.

extended trading time. After adjusting for the possible trade losses from other exchanges, we find that trading activity increases after the early opening.[5] This finding is important, given the recent extensions in trading hours (stocks can now be traded until 6:30 p.m.) and the move towards twenty–four hour trading. If longer trading time creates extra volume, then this would generate additional revenue which would, at least partially, offset the additional cost of operating for longer hours. Thus, the issue is interesting from both academic and policy perspectives.

Another important concern in extending daily trading hours is its effect on daily return variability. This is because risk premium is determined, at least in part, by the variability of asset returns. Various studies have found, within the appropriate windows, that longer trading hours do not increase return variability. Barclay et al. find that weekend (Friday – Monday) return variability increases with Saturday trading, but this increase is offset by decreases in Tuesday and Wednesday return variabilities. French and Roll (1986) find that there is no significant difference in weekly return variance with and without Wednesday trading. Booth and Chowdhury find that daily return variance does not increase after daily trading time was increased by an hour on the Frankfurt Stock Exchange. Consistent with these studies, we find that the extension of trading hours on the NYSE does not increase daily return variance.

The extension of trading hours on the NYSE also enables us to shed light on four documented empirical facts. This context provides alternative tests of the hypotheses proposed in the microstructure literature and, also, provides tests of some of the

[5] Barclay et al. find evidence that longer trading hours do not increase return variance *per se*, but they do not investigate the whether longer trading hours generates volume *per se*. However, if we assume that the increase in trade on the Tokyo Stock Exchange were not due to losses from other exchanges, their study would suggest that extending trading hours creates trade.

predictions of Hong and Wang (2000). The first two facts that we study are the patterns of intraday volume and return variability. It is well known that intraday volume pattern is U-shaped [see, for example, Jain and Joh (1988)] as is the return variability pattern [see, for example, Wood et al. (1985)]. Our goal here is to examine the effect of extending trading time on these intraday patterns. Our results show that both the volume pattern and the return variability pattern become more U-shaped at the open (thirty-minute pattern of the first two hours of trade) and at the close (thirty-minute pattern of the last two hours of trade).

Third, it is well established that returns are more volatile when markets are open than when they are closed. This phenomenon has been attributed to private information based trading [French and Roll (1986) and Barclay et al. (1990)], and public information arrival [Booth and Chowhury (1996)] with little evidence in favour of noise trading. The early opening provides us with another opportunity to shed light on the public information versus private information debate. This is because the early opening provides additional trading opportunity and, therefore, changes the arrival of informed traders at the open. If the private information hypothesis is true, these changes will be reflected in changes in the intraday return variances around the early opening. On the other hand, it is unlikely that the early opening will change the timing of public information releases around the new trading period. Thus, if the public information hypothesis is true, we do not expect changes in the intraday return variances around the opening. We find that intraday return variabilities around the new trading period change after the early opening in a manner consistent with the private information hypothesis.

Fourth, Amihud and Mendelson (1987) and Stoll and Whaley (1990) document that open-to-open returns display greater volatility than close-to-close returns. The early opening of the NYSE gives us the opportunity to examine the effect of extending trading time on transitory volatility in open-to-open returns relative to that of close-to-close returns. We find that extending trading time reduces transitory volatility in open-to-open returns relative to the close-to-close return transitory volatility. This evidence is consistent with the price formation hypothesis and confirms the findings of Gerety and Mulherin (1994).

The rest of the chapter is organized as follows. Section 2.2 discusses the various theories and their predictions regarding the effects of extending trading hours. Section 2.3 presents a brief discussion of the sources of the data. Section 2.4 reports the results, and section 2.5 concludes the chapter.

2.2. REVIEW OF THEORIES AND TESTABLE HYPOTHESES

There are basically two types of models that offer predictions regarding the effects of extending trading time on volume and return volatility. They are the strategic trader models [eg. Kyle (1985), Admati and Pfleiderer (1988), and Foster and Viswanathan (1990)] and the competitive trader models such as Hong and Wang (2000). Whereas the strategic trader models rely on the strategic behaviour of investors to derive the equilibrium outcome, the competitive trader types rely on the competitive behaviour of investors to reach equilibrium. The strategic trader models do not specifically address the effects of extending trading hours, but by carefully considering the microstructure of trading among informed traders, market markers, and liquidity traders, one can extract predictions of the effects of extending trading time on volume and return variability. Hong and Wang (2000) develop the first model that addresses the effects of extending trading hours on daily volume and return volatility, as well as on some of the observed empirical facts. We discuss the relevant predictions of these models and some of the testable microstructure hypotheses in the remainder of section 2.2.

2.2.1. The Effects of the Early Opening on Daily Volume and Return Variability

In Kyle's (1985) model, there are three types of traders: informed traders who trade strategically to maximize profits from their private information, liquidity traders who buy and sell randomly, and a market marker who sets the price based on the total sell and buy orders. The specialist cannot distinguish between liquidity and informed traders' orders and, hence, faces a signal extraction problem in setting prices to reflect private information. Within this setting, private information is incorporated into prices over time at a constant rate per trading hour. The price of the relevant trading interval reflects all private information and the variance of the return over the interval only reflects new information. Price and return variance are, however, influenced by random liquidity trade within the interval because of the signal extraction problem.

Admati and Pfleiderer (1988) and Foster and Viswanathan (1990) extend Kyle's strategic type model to include discretionary liquidity traders. These traders, unlike the random liquidity traders, choose when to trade. Thus, if discretionary liquidity traders believe that the informed traders have superior information, they may hold back their trades to reduce the likelihood of bearing the cost of trading with an informed trader. Admati and Pfleiderer (1988) show that a pooling equilibrium minimizes the trading cost of discretionary liquidity traders and hence, in equilibrium, they will choose to trade at the same time.[6] This pooling of discretionary liquidity trades attracts informed traders since this is the best time to disguise their trades. In Foster and Viswanathan (1990), private information at the beginning of each period becomes less valuable through the period because of the public announcement of some portion of it. Informed traders, therefore, trade more aggressively at the open in their setup.

In these strategic type models, the effect of extending daily trading time will critically depend on how private information production and volume of liquidity trade relate to the length of trading time. If the volume of liquidity trade remains unchanged because the needs for liquidity arise outside the markets [Admati and Pfleiferer (1988)], the longer trading time will result in thinner random liquidity trade per period (per 30-minutes for instance).[7] This will make masking informed trade more difficult and, hence, informed trade will be less per period. The reason is that a fall in the variance of liquidity trade per period will reduce market depth, which will discourage insider trading per period. Given the amount of private information, however, daily informed trade will remain unchanged as informed traders proportionately reduce their trades per period. This, together with unchanged daily volume of liquidity trade, will result in unchanged daily *total* volume. Thus, the models predict that extending trading time will not increase daily volume for a given private information production.

As discussed previously, the price of the relevant trading interval reflects all private information and the variance of the returns over the interval only reflects new information in these models. Thus, to the extent that an increase in trading hours does not

[6] The actual timing of concentration of trades and price changes are not determined in the models.

increase daily information production, the models predict that daily return variability will remain the same.

Intraday return variances would, however, be influenced by random liquidity trade within the intervals because of the signal extraction problem of the market maker. The models, therefore, predict that the pattern of intraday return variance would change reflecting the effects of the extension in trading periods on the arrival of traders. In particular, liquidity needs that arise overnight and must be satisfied at the earliest opportunity will be shifted from 10:00 a.m. - 10:30 a.m. to the new trading period (9:30 a.m. – 10:00 a.m.). If volume in the new period is substantial, informed traders would shift some of their trades from the surrounding period to the new period. Also, if some discretionary liquidity traders shift their trades to the new period [Admati and Pfleiderer (1988)], the shift of informed trades will be reinforced. Thus, the models suggest that return variance between close to 10:00 a.m. will increase because of the additional private information revealed through trading during this period while intraday return variabilities in the periods immediately after 10:00 a.m. will decline. An increase in the close-10:00 a.m. return variance along with a decrease in the return variance of the subsequent periods will, therefore, be consistent with migration of informed trades from the surrounding periods to the new trading period.

In Hong and Wang (2000), trade in stocks is motivated by two factors: private information and hedging needs. In their framework, investors' portfolios consist of traded assets (a stock and a money market account) and private investments.[8] When returns to the private investments and the stock are correlated, an investor's stock holding will be

---

[7] With an unchanged daily volume of random liquidity trade and longer period over which to spread it, the volume of each random arrival will be less. This will reduce the variance of the random liquidity trades.

driven by both the expected return on the stock and the expected return on the private investment. Specifically, investors can use the stock to hedge the risks from their private investments. This gives rise to hedging trade in the model. Apart from this, some investors receive private information about the future stock payoffs and they take speculative positions to capture the benefits from their private information.

The effect of extending trading time on daily volume in this model will depend on how speculative and hedging trades are influenced by the extended trading period. Speculative trade in the model is motivated by a constant and continuous flow of exogenous private information about the future payoffs of the stock to the class of informed investors. Since information flow is exogenous, the extension will not change the amount of new information. The early opening shortens the closure period and, consequently, reduces the non-trading risk associated with holding stocks overnight. The reduction in risk will increase speculative trade near the close since it is less costly to take speculative positions. On the other hand, the reduction in non-trading risk will decrease hedging trade near the close and subsequently at the open. Thus, the model predicts an increase in speculative trade and a decrease in hedging trade after the early opening. The overall effect on daily volume will, therefore, depend on whether the speculative trade dominates the hedging trade. An increase in daily volume after the early opening will be consistent with a dominant speculative motive, and a decrease in daily volume will be consistent with a dominant hedging motive.

In the Hong and Wang model, the equilibrium price of the stock is determined by private information, which gives rise to speculative trades, and private investment technological shocks, which gives rise to hedging trades. Since both private information

[8] Private investments can be viewed as illiquid assets.

and private technological shocks are exogenously determined, the early opening will not exert any systematic influence on daily return variability. Thus, the model predicts that daily return variability will remain the same after the early opening.

2.2.2. The Effects of the Early Opening on Some Empirical Facts

We also examine the effects of the early opening on four empirical observations in this study. These observations include the following.

(i) intraday volume exhibits a U-shaped pattern.
(ii) intraday return variability exhibits a U-shaped pattern.
(iii) hourly returns are more volatile during trading time than non-trading time.
(iv) open-to-open returns have more transitory volatility than do close-to-close returns.

These empirical facts are robust with respect to different market microstructures such as the NYSE, Nasdaq, and the interbank market of currencies. The objective of analyzing the effects of the early opening on these empirical regularities is threefold. First, this exercise will help us predict the effect of continuous trading on these observations, which has become important in view of the current move towards continuous trading. Second, analyzing the effects of the early opening on observations (iii) and (iv) provides alternative tests of some of the hypotheses proposed in the market microstructure literature and, hence, sheds light on these hypotheses. Third, examining the effects of the early opening on these empirical observations constitutes tests of some of the predictions of Hong and Wang (2000). The model by Hong and Wang is the first to capture these phenomena and to make various predictions of the effects of extending trading time on these observations.

2.2.2.1. Intraday volume and return exhibit U-shaped patterns

It has been observed that the intraday volume pattern is U-shaped [see, for instance, Jain and Ord (1988) and Chan et al. (1996)]. However, the influence of trading time on this pattern has not been investigated. Among others, [see, for example, Brock and Kleidon (1992)], Hong and Wang (2000) use market closures to derive the observed intraday volume pattern. In their model, closure increases the cost of holding the stock because it introduces non-trading risk. This leads investors to liquidate some of their hedging positions near the close of trade resulting in high closing volume. At the open, non-trading risk is eliminated and investors gradually re-establish their hedging positions resulting in high opening volume. Furthermore, speculative trade at the open is high due to accumulated overnight information. Thus, market closure, by introducing liquidity risk, results in high hedging trade at the close and the open and, by precluding investors from trading, results in high opening speculative trade. The high volumes at the open and the close give rise to the observed U-shaped intraday volume pattern in their model.

An increase in trading time (a reduction in the length of closure) will reduce non-trading risk and lead to less hedging trades at the close and at the open. On the other hand, while the shorter non-trading period will decrease speculative trade at the open (less accumulated overnight information), it would increase speculative trade near the close due to the reduction in the risk of holding overnight speculative positions (that is, for a given private information, informed traders will take more speculative positions near the close after the early opening). Consequently, the model predicts that speculative and hedging trades will fall at the open, but at the close, hedging trade will fall and speculative trade will increase. Thus, the model predicts that intraday volume pattern will unambiguously become less U-shaped at the open. At the close, the effect will depend on whether the hedging trade effect dominates the speculative trade effect. A dominant hedging trade will result in a less U-shaped pattern at the close, while a dominant speculative trade will result in a more U-shaped pattern at the close. Intuition, however, suggests that the hedging trade effect will dominate. The reason for this is that in a continuous trade setting, the level of intraday trading activity will be the same because information flows about stock returns and investors' private investments are exogenous and homogenous over trading and non-trading time. Thus, a reduction in closure will reduce the variation in intraday volume and, hence, the model predicts a less U-shaped volume pattern at the close.

It is also well documented that the intraday return variability pattern is U-shaped [see, for example, Wood, McInish, and Ord (1985) and Kleidon and Werner (1996)]. In Hong and Wang, this phenomenon is explained by intraday variation in the level of information asymmetry and hedging trade. Information asymmetry is high at the open (due to closure) and decreases as trading progresses through the day. Consequently, the stock price becomes more volatile through the day as it becomes more sensitive to information about its future payoffs (more responsive to private information arrival). On the other hand, overnight non-trading risk reduces hedging demand at the close causing the stock price to be less sensitive to investors' technological shocks. Consequently, stock return volatility decreases during the day. The U-shaped intraday pattern of return variability is obtained when the time varying hedging demand dominates early in the day and the effect of a decrease in information asymmetry dominates near the close.[9]

An increase in trading time in this setup will reduce non-trading risk, which will decrease the time variation in hedging demand, and also reduce information asymmetry. Thus, return variability will be less U-shaped at the open due to less time varying hedging demand and, also, at the close due to less time varying information asymmetry. Intuitively, closure gives rise to time variations in information asymmetry and hedging trade which generate the U-shaped return volatility. A decrease in closure will, therefore, reduce the intraday variations in information asymmetry and hedging trade and, consequently, the intraday variation in return volatility will decrease. Thus, the model predicts that intraday return variance will become less U-shaped after the early opening.

2.2.2.2. Hourly returns are more volatile during trading time than non-trading time

Fama (1965) and Granger and Morgenstern (1970) have documented that returns are more volatile during trading periods than non-trading periods. The reasons for this phenomenon have also been studied closely. French and Roll (1986) and Barclay et al. (1990) find that this phenomenon is consistent with private information based trading, and suggest a limited role for noise trading. By contrast, Booth and Chowdhury (1996) find that, in Germany, the phenomenon can be explained by public information. The early opening of the NYSE offers the opportunity to test whether the higher trading time return variance is consistent with private or public information.

---

[9] Although, Hong and Wang obtain a U-shaped pattern of return volatility for a wide range of parameter values, their model can produce monotonically decreasing and increasing patterns, as well as an inverted U-shaped pattern by changing the underlying parameters.

We have no reason to believe that the extension of trading hours on the NYSE will change the timing of public information releases in the morning. That is, there is no reason why public information releases will be moved from after 10:00 a.m. to the new trading period. This suggests that if the higher trading time return variability is due to public information releases, there should be no changes in the intraday return variances after 10:00 a.m.

On the other hand, if the higher trading time return variability is due to private information based trading, then changes in intraday return variances after 10:00 a.m. are to be expected. There are two major reasons for this. First, informed investors who desire to trade on their overnight information at the open to avoid decays will shift their trades from after 10:00 a.m. to the new trading period. Second, trades motivated by liquidity needs that arise overnight and must be satisfied at the earliest opportunity will be shifted from after 10:00 a.m. to the new trading period. Migration of liquidity traders will in turn attract some informed traders since the presence of liquidity traders helps disguise informed trades. In addition, if the volume of trade in the new period is high, it will attract some discretionary liquidity traders and this will reinforce the shift in informed trades to the new period [Admati and Pfleiderer (1988) and Foster and Viswanathan (1990)]. These shifts in informed trades will increase the return variance between the close and 10:00 a.m. after the early opening, and reduce the return variance in the surrounding periods. An increase in the close to 10:00 a.m. return variance, together with a decrease in intraday return variances immediately after the new trading period, will be consistent with the private information hypothesis.[10]

[10] This is similar to the Barclay et al. test. However, our tests can be differentiated on two accounts. First, our relevant trading period is daily, which is more appropriate for making inferences about daily extensions

Examination of the effect of the early opening on trading and non-trading return variabilities also provides a test of the prediction of the effect of extending trading time on these return variabilities in the Hong and Wang model. In this model, the returns over the trading periods are more volatile than the returns over the non-trading periods. Given a constant exogenous information flow and the same length of trading and non-trading periods, the model generates a higher volatility of stock returns during the trading period. This is because trading among investors reveals investors' private information which is incorporated into prices. The model, therefore, predicts that if the early opening does not increase the amount of new private information, then trading time return variability will not change relative to non-trading time return variability after adjusting for the effects public information, which works to increase trading time return variance.

2.2.2.3. Open-to-open returns have more transitory volatility than close-to-close returns

Open-to-open return variance is higher than close-to-close return variance [see, for example, Amihud and Mendelson (1987) and Stoll and Whaley (1990)]. Two hypotheses have been proposed for this phenomenon; the trading mechanism and the price formation. The trading mechanism arguments such as the use of call auctions [Amihud and Mendelson (1987)] and the participation of specialist [Stoll and Whaley (1990)] suggest that these mechanisms are responsible for the greater transitory noise at the open. The price formation argument suggests that overnight interruption of trade clouds prices and results in noisier opening prices [Dow and Gorton (1993), Leach and Madhavan (1993), and Romer (1993)]. Gerety and Mulherin (1994) use the Dow Jones

---

in trading periods. Second, Tokyo Stock Exchange and the NYSE likely have different characteristics that may influence the effects of longer trading. For example, institutional investors play a more active role on

65 returns from 1952 to 1992 to test these competing hypotheses. They find that the 24-hour return variances (open-open, 11:00 a.m.-11:00 a.m., ..., close-close), in general, decease through the day. Their evidence is consistent with the price formation hypothesis, which suggests that transitory noise in opening price would be reversed gradually through trading, and inconsistent with the trading mechanism hypothesis, which suggests an abrupt decline in the 24-hour return variance after the open.

The early opening of the NYSE provides an alternative means of testing these hypotheses. If the price formation hypothesis is true, the extension of trading hours on the NYSE would influence open-to-open return variability relative to close-to-close return variability in two ways. First, the decrease in the length of the overnight period will result in less noise at the open (less cloudiness), which reduces the ratio of open-to-open return variance to close-to-close return variance.[11] Second, the longer trading period would reverse more of the noise in opening price by the close of trade.[12] This will reduce the close-to-close return variance and reduce the ratio of open-to-open return variance to close-to-close return variance. Although there are reasons to believe the early opening will decrease both the open-to-open return variance and the close-to-close return variance, the price formation hypothesis suggests the former effect would dominate. This is because the price formation explanation hinges on market closure which clouds prices. A reduction in market closure will, therefore, reduce the disparity in the noise in opening prices and that in closing prices. On the other hand, there should be no change in

---

the Tokyo Stock Exchange and the trading mechanisms are different.

[11] It may be argued that the proportionate reduction in the overnight period is too small to induce a significant reduction in noise. However, it is likely that the generation of noise is concentrated during normal business hours when information flow is high (see, for instance, Hertzel et al. (1990)]. Since most firms are open for normal operations between 9:30 a.m. – 10:00 a.m., information flow is likely to be high relative the remaining closure period. This, coupled with closed markets, can generate disproportionately more noise during this period.

transitory volatility in opening prices relative to that in closing prices if the trading mechanism hypothesis is true. This is because the trading mechanism at the open did not change on the NYSE during the period under investigation. This suggests that a reduction in open-to-open return variance relative to close-to-close return variance after the early opening will be consistent with the price formation hypothesis and inconsistent with the trading mechanism argument.

Hong and Wang's model has implications for the ratio of open-to-open return variance and close-to-close return variance. In their model, the high level of information asymmetry at the open drives noise in open-to-open return variability.[13] As trade progresses during the day uninformed traders infer private information from the market price, and this reduces information asymmetry and the noise associated with it. As closing time approaches, overnight liquidity risk decreases hedging demand and this makes the stock price less informative about investors' private information. Thus, overnight liquidity risk, by inducing hedging trade near the close, introduces noise in the closing price. Transitory noise in open-to-open returns is higher than that in close-to-close returns when the noise in the opening price due to information asymmetry is higher than the noise in closing prices due to the decrease in hedging demand. The early opening decreases information asymmetry at the open, which decreases transitory noise in opening prices. It also reduces overnight liquidity risk, which slows down the decrease in hedging demand near the close and, consequently, reduces the transitory noise in closing prices as well. Thus, the effect of the early opening on the ratio of open-to-open return variance to close-to-close return variance is ambiguous in the model. The ratio will fall if

---

[12] Stoll and Whaley (1990), for instance, find that closing prices are *not* totally devoid of noise.

the effects of the reduction in information asymmetry dominate the effects of the decrease in hedging trade.

In summary, we test the effects of extending trading time on a number of market variables. These variables include daily volume and return variability, intraday volume and return variability, trading time return and non-trading time return variabilities, and open-to-open return variability relative to close-to-close return variability. The results of these tests help us to reach conclusions regarding the effects of extending trading time on volume and price behaviour, which has become important in light of the current move towards extending trading hours. They also provide evidence on the hypotheses proposed in the microstructure literature as well as test some of the predictions of Hong and Wang' model.

2.3. DATA DESCRIPTION

We use data from September 30, 1983 to September 29, 1987 as the core data for analyzing the effects of the early opening on the identified market variables. This represents two years of data before and after the NYSE started opening at 9:30 a.m. on September 30, 1985. We avoid using longer windows to minimize the effects of other extraneous factors (e.g. trends).[14] In addition, we use data from July 1 to December 31 for each of the years 1980 to 1990 to check the robustness of the results. This allows us to compare trading activity and return variances for the periods July 1 to September 29 and

---

[13] Noise in the sense that the market price deviates from the price which rationally assesses private information.

[14] Booth and Chowdhury (1996) use one year of data around January 15, 1990 (the day Frankfurt Stock Exchange started opening at 10.30 a.m. instead of 11.30 a.m. Frankfurt time). Using one year of data around the early opening does not change our conclusions.

September 30 to December 31 in 1985, and replicate these comparisons for "control" years 1980-1984 and 1986-1990.

The Standard & Poor's (S&P) 500 index return is used as the main market return, with the Dow Jones 65 (DJ 65) and the NYSE value-weighted market indexes used as supplementary market returns. We focus on the S&P 500 index returns because they are less likely to suffer from non-synchronous trading as well as capture, in a broad sense, the influence of the early opening on the market returns. Scholes and Williams (1977) and Dimson (1979) have shown that nonsynchronous and infrequent trading can induce autocorrelation in market indexes even when the true returns are not autocorrelated. This problem is likely to be especially severe in intraday market return studies. For this reason, the S&P 500 has been used in many studies to capture market returns [see, for example, Jain and Joh (1988)] and we adopt it in this study. We generate daily data for returns on the S&P 500 and the NYSE value weighted indexes from the Center of Research in Security Prices (CRSP) database, and DJ 65 intraday returns are obtained from Gerety and Mulherin.[15]

Daily data on volume of shares traded and the number of outstanding shares are also generated from the CRSP database. The number of firms with available data from CRSP for our period of interest (Sept. 30, 1983 – Sept 29, 1987) is 8,915. We filter out firms for which available data starts after the early opening because there are no bases for comparison for these firms. Apart from this, we remove firms that do not have data at least one year before the early opening and one year after it. This is to ensure that our study of the effects of the early opening is not significantly affected by inclusions and

[15] We wish to thank Mason Gerety and Harold Mulherin for supplying us with these data.

exclusions of firms from the CRSP database around the event date. After this filtering we have 2,676 firms, which forms the sample of firms for which we collect our daily data.

We also generate the volume of trade each thirty minutes from the Institute for the Study of Security Markets (ISSM) database. This results in twelve daily intervals for which volume is generated before the early opening and thirteen intervals the after the early opening. The ISSM database has 988 firms for which data are available at least one year before the early opening and one year after it. Among these are 19 firms for which data are not available on CRSP and we remove these 19 firms from the sample. This is because we use information on outstanding shares from the CRSP database to calculate intraday turnover for each firm. We calculate market turnover as the average of the turnovers of the individual firms.[16] Turnover, rather than volume, is used to gauge the level of trading activity in the market because turnover controls for new share issues and is less subject to trend than volume.[17]

2.4. EMPIRICAL RESULTS: THE EARLY OPENING AND TRADE

2.4.1. Regression Analysis

We calculate volume in the new trading period to be 11.6% of daily trade for the period September 30, 1985 to September 29, 1987.[18] We study whether some of this volume is due to the early opening in two ways. First, we examine the effect of the early opening on daily volume by using regression techniques and difference in mean tests. An increase in daily volume, which is not accounted for by shifts in trades from the other

[16] A more conventional measure of market turnover is total market volume divided by total market shares. This, however, assigns greater weight to the more actively traded stocks so we use the mean turnover across individual firms.

[17] Among others, Jain and Joh (1988) use turnover rather volume to adjust for the trend in volume.

exchanges (or trend or other seasonal factors), would be due to the longer trading period. Second, we account for the sources of the volume in the new trading period. In this case, volume in the new trading period could be due to shifts from other periods, shifts from other exchanges or trade induced by the longer trading time. This suggests that residual volume (volume after accounting for shifts from 10:00 a.m.- 4:00 p.m. and the other exchanges) would be due to the longer trading period.

Our study of the effects of the early opening on trading activity is based on its influence on daily market turnover, calculated as the mean daily turnover of the 2,676 firms on which at least one year of data is available pre- and post-early opening from the CRSP database. We use two years of data pre- and two years of data post-early opening on these firms to run the regressions.[19] The regressions consist of daily market turnover (TRN) as the dependent variable and the regressors are made up of lagged TRN (to correct for autoregressive processes - TRNi denotes the i-day lag), a dummy for the early opening, daily dummies (to remove day of the week effects), monthly dummies (to correct for monthly effects), a trend, and a trend squared. We experiment with twenty lags of the dependent variable, and remove the insignificant ones to obtain a more parsimonious equation. This specification does not affect the estimates of the coefficients of interest or their significance. Due to the time series nature of the data and the possible influence of trend, we run two regressions: one with trend, and the other with trend and its square. The results of these regressions are reported in Table 2.1. From the regression results, trend is significant at the 5% level in both regressions. It can be observed that the

[18] This figure is based on the sample of firms from the ISSM database.

[19] We also run regressions using data from 1980 to 1990. The results from these regressions do not change our qualitative results. Indeed, they offer stronger evidence in favour of an increase in trade after the early

inclusion of trend squared in the regression has little effect on the coefficient estimates other than on the trend. This suggests that trend has minimal effect on the coefficient estimates. In view of this, we focus on analyzing the effect of the early opening on volume by studying the results from the regression without the trend squared.

The results show that many of the variables included in the regression are significant at 5% level of test. There is a marked day of the week effect on volume, which is consistent with the finding in many other studies [see, for example, Jain and Joh (1988)]. The monthly effects are also significantly different from January with the exception of February, March and December.

Our variable of interest is the DUMMY variable and it takes a value of zero before the early opening and one after it. The coefficient estimate on this variable is positive and significant at 1% level of test in both regressions. The estimate of this coefficient is robust to the inclusion of the trend squared (it is 0.156 without the trend square and 0.163 with it). This suggests that the regression effectively controls for trend and, therefore, the DUMMY coefficient estimate is not unduly influenced by trend. The coefficient estimate on the DUMMY variable in the regression without the trend squared suggests that daily turnover increases by 6.3% after the early opening.[20]

To check the robustness of the observed increase in turnover after the early opening, we also use t-tests to compare the mean of daily market turnover three months before September 30 (1 July to September 29) and three months after September 29 (September 30 to December 31) for the years 1980 to 1990. The three-month windows

---

opening. We do not report these regressions because a longer window is more subject to the effects of exogenous shifts.

before and after the early opening are used to further minimize the effects of trend and other extraneous factors, and the years 1980-1984 and 1986-1990 serve as the control years, controlling for seasonal effects.

From the results (see Table 2.2), daily turnover increased by 19% in 1985 and this is significant at .01% level of test. This exceeds the increases in every other year apart from 1982. The increase in 1985, therefore, has an empirical probability value of 10%. In addition, while the increase in 1985 is robust to the length of the window, that of 1982 is not. It decreases from 31% for the three-month window to 27% and 20% for the two- and the one- month windows, respectively. This suggests that market turnover during the second half of 1982 is unduly influenced by trend. The trend during this period reflects the increase in volume which started with the inception of the bull market on August 17, 1982.[21] If we downplay the 1982 result as anomalous, or even otherwise, the evidence from the control years overwhelming suggests that the significant increase in daily volume three months after the early opening is, at least partially, due to the longer hours that the NYSE is open. In other words, both the regression analysis and the difference in mean tests suggest that daily turnover increases after trading hours were extended on the NYSE on September 30, 1985. The implication of these results is either trades are shifted from the other exchanges or the new trading period induces trade.

It is unlikely that the early opening of the NYSE attracts trades from the other American exchanges or the Canadian markets because they also opened at 9:30 a.m. effective September 30, 1985, to avoid possible losses to NYSE. Also, the new NYSE

[20] This is obtained by dividing the coefficient estimate on the dummy variable in the regression without trend squared by the mean turnover before the early opening, 2.49. Regression using data on all firms in the CRSP database suggests that daily turnover increases by 5% after the early opening.

trading period corresponds to an inactive period in Asia and, therefore, trade losses induced by the new trading opportunity may not originate from the Asian markets.[22] This leaves the European markets as the most likely location from which trades could be shifted to the NYSE. Indeed, institutional investors in Europe pressed for the extension to accommodate their continually increasing trading volume and interest in U.S. stocks (the Globe and Mail, July 11, 1985). These investors would shift their trades in US stocks from the European exchanges to the NYSE to take advantage of the lower transaction costs if the markets are simultaneously open.[23]

We use information on stocks listed on both the NYSE and the London Stock Exchange (LSE) to gauge the increase that may have been due to shifts in trade from the other exchanges.[24] If the increase in turnover observed after the early opening is due to losses from Europe, the stocks listed on both exchanges will display a significant increase in turnover relative to the stocks that were not listed on the LSE. From the regression analysis (see results in Table 2.3), the turnover of the stocks listed on both the LSE and the NYSE increases by 6.2%.[25] This is the same as the increase in turnover of the stocks listed on the NYSE but not the LSE (6.2%).[26] There is, therefore, no evidence that the observed increase in turnover on the NYSE after the early opening is due to volume losses from the other exchanges. This leads us to conclude that the increase in daily

---

[21] Jain and Joh (1988) also find a significant increase in turnover around August 17, 1982 and attribute this to the beginning of bull market.

[22] The new trading period on the NYSE corresponds to the late night hours in Asia when business activity is low and the markets are closed.

[23] As noted by Barclay et. al., for internationally listed stocks, transaction costs are likely to be lower on the domestic market than the foreign market.

[24] We use the LSE listing as a proxy for European listing, since the LSE is the largest market in Europe and major U.S. firms are likely to list there before listing elsewhere in Europe.

[25] The percentage increase is obtained by dividing the coefficient of the DUMMY variable by the average daily turnover of the dually listed stocks two years before the early opening, 2.60.

[26] The percentage increase is obtained by dividing the coefficient of the DUMMY variable by the average daily turnover of stocks not listed on the LSE two years before the early opening, 2.49.

turnover cannot be explained by trade losses from the other exchanges and, hence, the observed increase in trading activity must be due, at least partially, to the longer trading period.

2.4.2. Accounting for Volume in the New Period

Trading in the new period could arise because of trade migration from 10:00 a.m. - 4:00 p.m. (other periods) or from the other exchanges, or because the longer trading period generates trade. Since we estimate that there are no trade losses from the other exchanges in the previous section, we focus on determining the shifts from the other periods in this section.

To find the periods when the shifts occur and, also, their magnitudes, we use thirty-minute intraday data from the ISSM database. For each thirty-minute trading period (we refer to the thirty-minute trading periods as "sub-periods"), we calculate market turnover as the mean of the firms' turnovers in the sub-period. We then calculate the average of the market turnover for each sub-period two years before the early opening and two years after it. Plots of the average market turnover for the sub-periods show the intraday volume patterns two years preceding the early opening and two years after it.

With the exception of 10:00 a.m. – 10:30 a.m. sub-period, the average of the post-early opening turnovers are higher than their pre-early opening counterparts for all the sub-periods. This indicates a trend in turnover which is consistent with the finding of a significant trend coefficient in the regressions. Therefore, to identify the sub-periods where the shifts occur, we scale down the average post-early opening turnovers such that the pre- and the post-early opening mean turnovers are equal during the 11:30 a.m. - 12 noon interval (this essentially detrends the post-early opening mean turnovers if the daily trend is appropriately captured by the trend in turnover between 11:30 a.m. and noon). This task is accomplished by scaling down the post-early opening mean turnover by the difference between the pre- and the post-early opening mean turnovers between 11:30 a.m. and 12 noon. The displacement preserves the curvature of the post-early opening curve and helps identify periods when shifts may occur.

The scaled post-early opening and the pre-early opening mean turnovers are plotted in Fig. 2.1, and below them is the plot of their differences. Negative values in the "difference" plot identify times when shifts occur. The "difference" plot indicates that some trades are accelerated from 10:00 a.m. - 11:30 a.m. to the new trading period. We conveniently divide the trading period (10:00 a.m. – 4:00 p.m.) into four one-and-half-hour trading periods and calculate the percentage changes in turnover for these periods after the early opening (see Table 2.4). While the *raw* turnover decreases between 10:00 a.m. and 11:30 a.m., it increases for the remaining one-and-half hour intervals. This suggests that some trades are accelerated from 10:00 a.m. - 11:30 a.m. to the new trading period.

The shifts in trade could be due to acceleration of liquidity trade and/or noise trade [as defined by Black (1986)] and/or informed trade to the new trading period.[27] We test whether informed traders accelerate some of their trades by examining the changes in the sub-period return variances after the early opening. Acceleration of informed trades can be motivated by the acceleration of liquidity trades [Admati and Pfleiderer (1988) and Foster and Viswanathan (1990)] or by the acceleration of noise trades [Kyle (1985)]. We can distinguish between the acceleration of noise trades and the acceleration of discretionary liquidity trades by examining the changes in correlations among the sub-period returns after the early opening.

To test whether informed trades are accelerated from 10:00 a.m. – 11:30 a.m., we divide the trading period (10:00 a.m. – 4:00 p.m.) into four one-and-half-hour trading periods. We examine the effect of the early opening on the return variances in these sub-periods to ascertain whether some informed traders migrate from these periods. Table 2.5 shows the variance and relative (relative to close-to-close) variance of the S&P 500 returns in these sub-periods. We use relative variance in the analysis because the return variances increase in all the sub-periods.[28] The results show that the largest fall in relative return variance (11%) occurs between 10:00 - 11:30 a.m., followed by a 4% fall between 1:00 - 2:30 p.m.[29] We use t-tests to determine the significance of the decrease in the return variances. The test is based on the assumption that the mean sub-period relative return variance over a year does not change significantly. The results show that only the decrease in relative variance between 10:00 a.m. - 11:30 a.m. is significant, and this is consistent with the acceleration of informed trades from 10:00 a.m. - 11:30 a.m. to the new trading period.[30]

We test whether the acceleration of informed trades is motivated by shifts in liquidity and/or noise trades. A shift in noise trades can be studied by examining the correlation among intraday returns between the close of trade and the early hours of

---

[27] Black (1986) describes noise traders as traders who trade on the wrong information.
[28] The overall increase in variance in 1986 and 1987 may partially reflect the market instability prior to the October crash. Booth and Chowdhury (1996) also use relative variance to control for possible influences of Germany's reunification and the Gulf war on return variances in 1990.
[29] The observed shift in trade between 1:30 p.m. – 3:00 p.m. is partly reflected in the lower return variance between 1:00 p.m. – 2:30 p.m.
[30] Barclay et al. (1990) argue, based on reductions in return variances on Tuesdays and Wednesdays, that informed trades were shifted from these days to Saturday when the Tokyo Stock Exchange opened for

trading. A shift in temporary noise trading between 10:00 a.m.-11:30 a.m. to the new trading period will, for example, increase the negative correlation or the positive correlation among intraday returns between the close and 10:30 a.m. This will be due to the additional noise impounded in prices by trading between 9:30 a.m. – 10:00 a.m. The return between the close and 10:30 a.m. can be written as:

$$r_{close,10:30} = r_{close,open} + r_{open,10:00} + r_{10:00,10:30}$$

where $r_{t,t+1}$ is the return between times t and t+1, and $r_{open,10:00}$ is zero before the early opening. Information about the correlation among the returns can be obtained by calculating equation (1).

$$VR_{10:30} = \frac{\sigma^2(r_{close,10:30})}{\sigma^2(r_{close,open}) + \sigma^2(r_{open,10:00}) + \sigma^2(r_{10:00,10:30})} \qquad (1)$$

where $\sigma^2(r_{t,t+1})$ is the variance of the return between period t and t+1. If the intraday returns between the close and 10:30 a.m. are uncorrelated, $VR_{10:30}$ will be equal to one (the subscript, 10:30, denotes the last time the price is observed to calculate VR). A negative correlation will be reflected in a value of $VR_{10:30}$ below one and a value above one indicates a positive correlation among the returns. If the shift in noise trading to the new period gets reversed by 10:30 a.m., $VR_{10:30}$ will decrease after the early opening. To study the effect of the early opening on noise trading near the open, we calculate VR_i for i=10:30, 11:00, 11:30, and 12 noon. The results are reported in Table 2.6. From the table, $VR_{10:30}$ is less after the early opening (for both the one- and the two-year windows around the early opening). This is consistent with temporary noise trading in the new trading period. $VR_{11:00}$ through $VR_{12:00}$ are, however, higher post-early opening. This suggests

Saturday trading. Also, Booth and Chowdhury also find that informed traders shift their trades from the other trading hours to the new trading period in Germany.

that temporary noise trading between the open and 11:00 a.m. through 12:00 noon did not increase after the early opening. These results are consistent with a shift in noise trading to the new trading period and 10:00 a.m. – 10:30 a.m.[31]

Also, random liquidity needs that arise overnight and must be satisfied at the earliest opportunity will be moved from 10:00 - 10:30 a.m. to the new trading period after the early opening. It is, therefore, not surprising that the biggest fall in variance occurs between 10:00 and 10:30 a.m. (the relative variance declines from 26% to 6%, see Table 2.7).[32] This is consistent with acceleration of liquidity trade. Thus, the evidence suggests that both noise and liquidity trades are shifted from the other periods.

A way to measure the amount of trade that is accelerated from 10:00 a.m.-11:30 a.m. is to add up the trade losses from these periods after scaling down the post-early opening mean turnovers.[33] Using this approach, we calculate that 20% of the trade in the new period can be accounted for by shifts in trade from 10:00 a.m. - 11:30 a.m. The unexplained trade in the new period could be due to shifts from the other exchanges or the longer trading period. Based on information on dually listed stocks on the LSE and the NYSE, we find in section 2.4.1.1. that there is no evidence of shifts in trades from the other exchanges. Thus, about 80% of trade in the new period cannot be accounted for by either acceleration of trades from the other period or by shifts from the other exchanges.

---

[31] The less noise in prices between the open and 11:00 a.m. through 12 noon after the early opening is consistent with a reduction in the noise in opening prices. This supports the argument that the higher noise in prices between the close and 10:30 a.m. is due to shifts in noise trading.

[32] Note that the variance in the first 30 minutes of trade also reflects public information released after close of trade. Berry and Howe (1994) find evidence that a substantial proportion of public information is released after close of trade. Whilst public information moves prices it may not affect volume.

[33] This assumes that the scaling appropriately adjusts for the trend (i.e. there are no shifts from 11:30 a.m.-noon). To the extent that there are some shifts in this period, the method will underestimate the amount shifted.

This evidence supports the finding from the regression analysis that the extension of trading time on the NYSE in 1985 generates additional trading.

2.4.3. Other Evidence

In this section, we study other evidence on the relationship between trading time and volume. In particular, we examine the evidence Barclay et al. use from the Tokyo Stock Exchange, weekly turnover in 1968 on the NYSE, and also the effects of the extension of trading time on the NYSE on October 1, 1974 (closing time was shifted from 3:30 p.m. to 4:00 p.m.). Evidence from Barclay et at. (1990) research indicates that on the Tokyo Stock Exchange the three-hour Saturday trading increases weekly volume by 21%. In their experimental environment, the issue of trend does not arise. This is because the exchange opened for trading on some weekends and closed on others from January 1973 to January 1989.[34] The larger weekly volume would, therefore, be due either to migration of trade from other exchanges or to the longer trading period. Although the other Asian markets are closed on Saturdays, which raises the possibility of some trade losses, it is unlikely that the volumes on the other Asian markets decrease by the gain on the Tokyo market. Thus, the evidence seems to indicate that the Saturday trading on the Tokyo Stock Exchange increases volume, at least partially.[35]

The Wednesday closings of the NYSE in the second half of 1968 to clear up backlogs of paperwork provides another opportunity to examine the effect of trading time on volume. The closure of the exchange on Wednesdays effectively reduces the trading

[34] The exchange opened on 585 Saturdays and was closed on 254 Saturday during this period.
[35] Barclay et al. observed the increase in weekly volume with the Saturday trading, but they do not track the source(s) of the additional volume.

period in the week from the normal five-day to four-day trading.[36] Since the length of weekly trading time was reduced in the second half of 1968, this period will be associated with reduced weekly turnover if trading time and trading activity are positively related. We can examine this by comparing the average weekly turnover in the second half of 1968 to the weekly turnover in the first half of the year. The average weekly turnover in the first half of 1968 is 14.14 and in the second half is 11.82.[37] The average weekly turnover in the first half of the years 1963-1967 and 1969-1973 is 8.5, and for the second half it is 7.6. Thus, the average weekly turnover in the second half of 1968 is 2.32 less than that of the first half, while for the ten years surrounding 1968, the average weekly turnover in the second half of the years is only 0.9 less than the average in the first half. This evidence suggests that the reduced number of trading days per week decreases weekly trading activity in the second half of 1968.

On October 1, 1974, the NYSE extended it trading time by shifting its closing period from 3:30 p.m. to 4:00 p.m. (we also refer to this as the late closure). This thirty-minute increase in daily trading time is the same as the time gained on September 30, 1985 when the early opening was introduced. There is, however, no intraday data covering the 1974 period from the ISSM database and, hence, a comprehensive study such as the one we undertake for the early opening cannot be replicated here. Using data from the CRSP database, we assess the impact of the late closure on turnover.

To do this, we use data on daily market turnover two years before the extension and two years after it (as before, we filter out firms that do not have data for at least one

[36] Information from these shorter weekly trading periods provides the core data that French and Roll (1985) used to examine the reasons for the higher trading time return variance.

year before and after October 1, 1974). Among other things such as minimizing the effect of trend, the choice of the window also facilitates comparison of the results with those of the early opening. We report the regression results for this period in Table 2.8. Although the DUMMY coefficient estimate suggests that turnover increases by 1.9%, the increase is not significantly different from zero.[38] It should be mentioned that t-test suggests that there is no significant increase in daily mean turnover three months after the late closure and regression analysis suggests that there is no significant increase after one year. Also, the late closure does not increase daily return variability. The F-statistics are 1.09 and 1.13 for one- and two-year windows around the late closure with significance probabilities of 0.495 and 0.164, respectively.

In contrast to our earlier result that the longer trading period increases the turnover after the early opening, the late closure does not increase trading activity. In the next section, we focus on determining the sources of the increase in turnover after the early opening to shed light on our empirical findings.

2.5. EXPLANATIONS FOR THE INCREASE IN TRADE POST-EARLY OPENING

The observed positive relationship between trading time and turnover after the early opening suggests that an increase in trading time can increase trading activity.[39] The result regarding the effects of the early opening is inconsistent with the predictions of the strategic trader models unless there is an increase in private information production or

---

[37] Average weekly turnover is calculated as the average of the weekly turnovers of the individual firms. Weekly turnover for each firm is obtained by dividing the weekly volume by the number of outstanding shares.

[38] The coefficient estimate is divided by average daily turnover before the last closure (1.0636).

noise trading [see, for instance, Kyle (1985)]. The evidence is consistent with a dominant speculative trade effect in the Hong and Wang's model. Their model suggests that the early opening can increase speculative trade because it reduces overnight risk (stronger reaction to given information arrival near the close) and reduces information cancellation (random arrival of information during closure). Based on the literature, therefore, the longer trading time can generate additional volume if it increases private information production, decreases information cancellation, increases noise trading, or increases trading at the close due to the lower overnight risk.[40]

The various possible explanations of the extra volume have different testable implications. An increase in private information production will be reflected in an increase in daily return variance after the early opening. A decrease in information cancellation, on the other hand, will increase the daily volume of trade but not daily return variability. If the extra trading activity post-early opening is due to an increase in noise trading that move prices, then we would expect prices to be noisier between the close and 11:30 a.m. after the early opening. This is because the extra noise trading has to occur between the open and 11:30 a.m. since there is no evidence of trade migrations from the periods after 11:30 a.m. to the new period. We can, therefore, test whether the data are consistent with an increase in noise trading by comparing the correlations among the intraday returns between the close and 11:30 a.m. An increase in negative correlations of the intraday returns between the close and 11:30 a.m. would be consistent with an

[39] Although the relationship between trading time and trading activity is not significant after the late closure, we press on with finding the sources of the observed increase in trading activity after the early opening which sheds some light on the observed insignificant increase in trade after the extension in 1974.
[40] Note that a reduction in overnight risk increases speculative trade at the close but decreases hedging trade. Thus, a higher closing volume will be consistent with a dominant speculative trade.

increase in noise trading.[41] A reduction in overnight risk will increase speculative trade near the close. We can, consequently, investigate the predictions of each of these sources to identify the reason(s) for the observed increase in volume.

2.5.1. Private Information Hypothesis

To determine whether the observed increase in volume is due to an increase in private information production, we compute the variances of daily close-to-close S&P 500, the DJ 65, and the NYSE value-weighted index returns two years before and two years after the early opening. The S&P 500 returns variance increases by 66%, the DJ 65 return variance increases by 30%, the NYSE value-weighted index return variance increases by 33%. The increases in the variances of the market returns are significant both by their F-statistics and the modified Levene W10 statistics.[42] For the three-month window around the early opening, the modified Levene W10 statistic and the F-statistic indicate that there is no significant increase in the return variances.[43]

Apart from calculating the variances of the returns of the market indexes, we also examine the distribution of the absolute values of the returns before and after the early opening as an additional test of the variability in returns. The results confirm the finding from the variance calculations. That is, there is no significant difference in the mean absolute returns of the market indexes three months before the early opening and the

[41] Note that a decrease in price reversals does not imply a decrease in noise trading. This is because of the possibility of an increase in smart trading (trades by investors who know the counter-parties are wrong) and, also, the reduction in the overnight period could reduce the cloudiness of opening prices.
[42] The Levene statistic does not depend on any distributional assumption and, consequently, avoids the assumption of normality which is used to derive the F-statistic.
[43] For the S&P 500 returns, for example, the F-statistic is 1.2 and the median Levene W10 statistic is 0.26 with significance probabilities 0.48 and 0.61, respectively.

three months after it. For the two-year window, the means absolute returns after the early opening is significantly higher than the means two years before.

Since the early opening has a significant effect on turnover within one month, it is unlikely that its effect would not manifest itself in higher daily return variability in three months. This argument is particularly valid if the private information hypothesis is true. This is because the private information hypothesis asserts that the observed increase in trading activity is driven by a higher level of private information production.[44] Thus, the significantly higher level of trading activity observed in the first three months is expected to be associated with a significant increase in daily return variability. This deduction from the private information hypothesis offers a more precise testable implication (a test that can, potentially, differentiate increases in return variance induced by informed trading from a general increase in return variability caused by other factors). Since private information can only be incorporated in prices during trading hours, the private information hypothesis predicts that the increase in return variability is due to an increase in trading time return variability and not non-trading time return variability. This suggests that trading time return variance will increase relative to non-trading time return variance after the early opening. We, therefore, examine the effect of the early opening on trading time return variability relative to non-trading time return variability to further assess the validity of the private information hypothesis.

We calculate non-trading time return as the value-weighted return between close and 10:15 a.m. pre-early opening and close - 9:45 a.m. post-early opening for the sample

---

[44] French and Roll (1986) point out that the private information hypothesis is consistent with a higher level of private information production during trading period and trading on private information that accumulates during the non-trading period. The former explanation suggests that if the increase in volume is due to private information, then trading time return variability will proportionately be higher than the non-trading

of our stocks from the ISSM database [since most stocks have opened within 15 minutes of the opening bell - see, for example, Stoll and Whaley (1990)]. For the two-year window around the early opening, the results show that the non-trading time return increases from 1.9×10^{-5} to 2.5×10^{-5} and the trading time return variance increases from 4×10^{-5} to 5.8×10^{-5} after the early opening. Thus, the percentage of the non-trading time return variance to the trading time return variance decreases from 47% to 43%.

The post-early opening trading time return variance includes the variance between 9:45 a.m. – 10:15 a.m. that was previously captured in the non-trading time return variance. This variance is latent but it can be estimated. We do this by noting that overnight return variance is largely explained by public information since private information and noise are mainly captured in returns during trading. From May 1990 to April 1991, approximately 6% of the overnight public information, as measured by news releases by Reuter's News Service, arrives between 9:45 a.m. and 10:15 a.m. [Berry and Howe (1994)]. This information suggests that the return variance between 9:45 a.m. - 10:15 a.m. is 6% of the overnight variance. After the adjustment for the shift in public information between 9:45 a.m. - 10:15 a.m. the variance of the non-trading time return relative to the trading time return variance did not change (47%).

The above results lead us to conclude that the observed increase in return variability for the two years window is not due to an increase in private information production. In view of the importance of this conclusion, we wish to summarize the compelling reasons for it.[45] First, the results from the three-month window suggest that there is no significant increase in return variability after the early opening although

time return variability. Information accumulation, on the other hand, cannot explain the observed increase in volume unless we consider the effect of information cancellation during the non-trading period.

turnover displays a significant increase in the first month. Second, the three-month window results are consistent with the findings of many other researchers [for example, French and Roll (1986) and Barclay et al. (1990) do not find evidence of an increase in weekly return variability for weeks with more trading days, and Booth and Chowdhury (1996) do not find evidence of an increase in daily return variance when trading time was extended by an hour on the Frankfurt Stock Exchange]. Third, and more importantly, the trading time return variance did not increase relative to non-trading time return variance after the early opening. Our results, together with the findings of other authors, support the conclusion that the early opening did not significantly increase return variability and that the observed increase over the two-year period is due to other factors. One possible factor is the instability that characterized the market prior to the October 1987 crash. Our evidence is, therefore, inconsistent with the prediction of the private information hypothesis.

2.5.2. Noise Trading Hypothesis

The noise trading hypothesis predicts that the early opening increases noise trading. If the new trading period creates additional noise that is temporary, we would expect increases in reversals among the intraday returns between the open and the surrounding periods. For example, if part (or all) the noise in the new period gets reversed in the next thirty minutes of trade, then there would be more reversals between the open and 10:30 a.m. after the early opening due to the additional reversals of the new trading period's noise. We, therefore, examine the importance of temporary noise trading in increasing the volume of trade after the early opening by studying the correlations

---

[45] It suggests that private information production is not related to trading time.

among the thirty-minute returns near the open.[46] The return between the close and 11:00 a.m. can be written as,

$$r_{close,11:00} = r_{close,open} + r_{open,10:00} + r_{10:00,10:30} + r_{10:30,11:00}$$

where $r_{t,t+1}$ is the return between time t and t+1, and $r_{open,10:00}$ is zero before the early opening. Information about the correlation among the returns can be obtained by calculating the variance ratio in equation (1). Additional temporary noise trading, which gets reversed by 11:00 a.m., will be captured by a decrease in equation (1) after the early opening. We, thus, study the change in temporary noise trading by referring to the calculations of $VR_{11:00}$, $VR_{11:30}$, and $VR_{12:00}$ in Table 2.6. From the results, the pre-early opening $VR_{10:30}$ is higher than the post-early opening $VR_{10:30}$, but the pre-early opening $VR_{11:00}$ through $VR_{12:00}$ are less than the corresponding post-early opening variance ratios. The results are, therefore, inconsistent with an increase in noise trading between the close and 12 noon. Thus, the evidence suggests that the increase in trading activity after the early opening is not due to an increase in noise trading.

2.5.3. The Lower Overnight Risk Hypothesis

We can gauge the increase in trade near the close by examining the ISSM intraday plots in Fig. 2.1. The plots represent the thirty minutes intraday volume patterns pre- and post-early opening. The post-early opening mean turnovers are higher than the pre-early opening ones (with the exception of the average turnover between 10:00 a.m. - 10:30 a.m.) and, hence, the post-early opening turnover for every period is scaled down such that the pre- and the post-early opening turnovers are equal at noon. If this scaling

[46] Note that this test will not capture noise trading that takes more than one day to reverse. Also, if the early opening increases smart trading (trade by investors who know their counter-parties are wrong), an increase

appropriately adjusts for trend, then the figure suggests that there is an increase in turnover in the last thirty minutes of trade after the early opening. From the figure, turnover in the last thirty minutes of trade increases by 0.04, and this represents 25% of the increase in daily turnover after the early opening. This estimate may, however, be exaggerated as the figure shows that some trades might have been shifted from 11:00 p.m. – 3:00 p.m. to the close.

It should be pointed out that the displacement of the post-early opening turnovers such that the pre- and the post-early opening turnovers are equal at noon is arbitrary. It is, therefore, useful to discuss the effect of scaling the post-early opening turnovers such that the equality occurs at other times apart from the periods immediately following the new trading period (10:00 a.m. - 11:00 a.m.). It is inappropriate to scale down the post-early opening intraday turnovers such that the pre- and the post-early opening turnovers equal at 10:30 a.m. or 11:00 a.m. because the data suggest that these early trading periods lose some trades to the new trading period and, hence, such displacements will underestimate the trend in turnover. In view of this, we ignore turnover between 10:00 a.m. – 11:00 a.m. in considering the possible trading times that the pre- and the post-early opening turnovers should equal to reasonably adjust for trend. With this exclusion, we calculate that the maximum increase in daily turnover that can be accounted for by increases in speculative trade near the close (3:30 p.m. – 4:00 p.m.) is 30%. This is obtained when the post-early opening turnover is displaced such the pre- and the post-early opening turnover is equal at 11:30 a.m. or 2:00 p.m. Thus, the maximum increase in turnover near the close is insufficient to account for the observed increase in trade.

---

in noise trading may not be reflected in price changes.

2.5.4. Information Cancellation Hypothesis

The information cancellation hypothesis provides a simple explanation for the observed relationship between trading time and turnover. This explanation hinges on the random nature of information arrival. Market closure prevents trading on information as it arrives. Since the arrival of good and bad news is random, it is logical that information cancellation (or partial cancellation) occurs when the market is closed. The amount of information cancellation that occurs in a particular period during market closure will depend on the accumulated information at the beginning of the period and the arrival of information during the period.

In general, more information will accumulate the longer the closure period. Thus, information accumulation at 9:30 a.m. will, on average, be higher than the earlier closure periods. Also, the sub-period 9:30 a.m. - 10:00 a.m. is economically significant in the sense that most firms have open for normal business and the rate of information arrival is generally high.[47] French and Roll (1986), for example, point out that information flow may be high during the normal business hours because information generating activities such as visiting corporate headquarters, examining company documents, and making recommendations to clients are all easier during this period. The fact that information has accumulated overnight (before 9:30 a.m.) and the high level of information arrival between 9:30 a.m. - 10:00 a.m. suggest that information cancellation will, on average, be higher between 9:30 a.m. and 10:00 a.m. than the other 30-minute closure periods. As a result, the new trading period after the early opening prevents significant information

[47] Also, see Hong and Wang (2000) where economic time is defined by the amount of information arrival. Hertzel et al. (1990) find that information flow about a currency is high during the business hours of the country that the currency originates.

cancellation, which explains (at least partially) the increase in trading activity after the early opening.

By contrast, extending trading time into insignificant economic periods may not generate significant trading activity as evidenced by the extension of trading time on October 1, 1974. Information cancellation is not likely to be high between 3:30 p.m. - 4:00 p.m. This is because, although, information arrival may be high during this time, accumulated information at 3:30 p.m. will, on average, be low since there is no closure before this time. Thus, it is unlikely that significant information offsetting occurred during this period when the market was closed. This suggests that the extension of trading time to 4:00 p.m. does not significantly preclude information cancellation, which is why a significant increase in trading activity is not visible. The changes in turnover after the extensions of trading time in 1974 and 1985 are, therefore, consistent with the information cancellation hypothesis.

Since this hypothesis has not been developed in any formal way, we present a brief model on the effects of extending trading time on information cancellation and, consequently, on the volume of trade, in the appendix. The model shows, under fairly weak assumptions about trade, that random information arrival during market closure results in less volume of trade than would otherwise be obtained if the market were not closed. Consequently, we show that extending trading hours generates additional trading activity.

2.6. EFFECT OF THE EARLY OPENING ON SOME EMPIRICAL OBSERVATIONS

2.6.1. Intraday Volume and Return Variability Patterns

We examine the effect of extending trading time on intraday turnover patterns by determining whether the post-early opening curves are more or are less U-shaped. We do this by focusing on the curvatures near the open and the close of trading. Since the post-early opening trading period is longer than the pre-early opening period, we drop turnover between 11:30 a.m. - noon in the post-early opening trades so that we can compare the curvatures near the open (the 30-minute trades in the first two hours of trade) and the close (the 30-minute trades in the last two hours of trade). The average 30-minute turnovers pre- and post-early opening (the average is taken two years before and two years after the early opening) are shown in Fig. 2.2, with their difference below it.[48] An (A) increase (decrease) in the 'difference' curve near the open or the close indicates that the post-early opening intraday pattern is more (less) U-shaped at the open or at the close.[49] The 'difference' curve increases at both the open and the close, but more so at the close. Thus, there is evidence that the intraday turnover pattern is more U-shaped after the early opening.[50] The increase in convexity following the early opening could be due to shifts in trade to the opening and to the closing periods. While the increase in convexity near the open is inconsistent with the prediction of Hong and Wang's model, the increase near the close is consistent with a dominant speculative trade effect.

---

[48] The graphs display the U-shape curve typically observed in intraday volume data [for example, Jain and Joh (1988) document U-shaped hourly aggregate volume using data from the NYSE from 1979 to 1983].

[49] We use data on the 30-minute turnover two hours after the market opens and the two hours before it closes to determine the effect of the extension on the curvature near the open and the close.

[50] The increase at the open is inconsistent with the predictions of Hong and Wang (2000). At the close, however, it is consistent with a dominant speculative trade effect.

The thirty-minute return variances of the S&P 500 two years before the early opening and two years after it are plotted in Fig. 2.3. The figure shows that intraday return variances become more U-shaped post-early opening. The increase in convexity is consistent with shifts in informed trades from 10:00 a.m. - 11:30 a.m. to the new period and from 1:30 p.m. - 3:00 p.m. to the last thirty minutes. Again the increase in convexity near the close is consistent with an increase in speculative trade near the close as informed traders react more aggressively to new information because of the lower overnight risk of speculative positions [Hong and Wang (2000)]. At the open, however, the increase in convexity of intraday return variance is not consistent with Hong and Wang's model.

2.6.2. Trading versus Non-trading Return Volatility

A well-established fact in market microstructure is that returns are more volatile when markets are opened than when they are closed [see, for example, Fama (1965) and Granger and Morgenstern (1970)]. The early opening gives us the opportunity to shed light on the private versus the public information debate.[51] We test this by calculating the changes in intraday return variances around the new period. From information on the returns on stocks from the ISSM database, return variance between 10:15 a.m. – 10:45 a.m. decreases from 7.4% to 4.9% after the early opening. On the other hand, the return variance from close - 10:15 a.m. as a percentage of close-to-close return variance increases from 34% to 41% after the early opening.

[51] We do not investigate the noise trading hypothesis since there is enough evidence in the literature that noise trading plays a trivial role in explaining the higher return variability during the trading periods. In addition, we find in section 2.5.2. that temporary noise trading did not increase after the early opening.

The decrease in 10:15 a.m. - 10:45 a.m. return variability and the increase in close - 10:15 a.m. return variance cannot be explained by the public information hypothesis unless public information releases were shifted from 10:15 a.m.-10:45 a.m. to the close - 10:15 a.m. period after the early opening. There is, however, no reason to believe that the early opening changed the timing of public information releases around 10:15 a.m. Therefore, it seems unlikely that the public information hypothesis can explain the observed changes in the return variance around 10:15 a.m.

The observed changes are, however, consistent with the private information hypothesis. The private information hypothesis asserts that the higher trading time return variability is due to private information which gets impounded in prices in course of trading. Thus, the private information hypothesis suggests that the increase in close - 10:15 a.m return variance and the reduction in 10:15 a.m. - 10:45 a.m. return variance are due to shifts in informed trades from 10:15 a.m. - 10:45 a.m. to the close – 10:15 a.m.[52] As noted earlier informed traders will accelerate their trades because of the acceleration in noise trades and in liquidity trades (liquidity needs that arise overnight and must be satisfied at the earliest opportunity). Our results are, therefore, consistent with the private information explanation of the higher trading time return variability.

[52] A number of reasons make this explanation plausible. First, traders would trade as early as possible on overnight information which might decays rapidly. Second, liquidity needs that arise overnight and must be satisfied at the earliest opportunity will be shifted to the new trading period, and this will attract informed traders [Admati and Pfleiderer (1988)].

2.6.3. Opening versus Closing Transitory Noise

We examine the effect of extending trading hours on transitory volatility in opening prices relative to closing prices by comparing the pre-early opening ratio of open-to-open return variance to close-to-close return variance to that of the post-early opening period.[53] This also sheds light on the two explanations offered for the observed higher transitory noise in opening prices than in the closing ones: the trading mechanism and the price formation hypotheses. While volatility due to noise (cloudiness of prices) would be reversed during trading time, the volatility associated with private information would be permanent. We can, therefore, determine the noise component in opening prices by checking for reversals in returns during the trading period.

$$\frac{\sigma^2(r_{o,t})}{\sigma^2(r_{c,t})} \qquad (2)$$

If the price formation hypothesis is true, two opposing forces will be exerted on equation (2). The shorter closure period will result in less noise at the open [this decreases equation (2)]. On the other hand, for a given noise at the open, the longer trading time will reduce more of the noise by the close of trade [this increases equation (2)].[54] If the trading mechanism hypothesis is true, equation (2) will not change since the NYSE did not change its trading mechanisms during the period under investigation. Thus, a change in equation (2) will provide support for the price formation hypothesis.

The ratio of open-to-open return variance to that of close-to-close, equation (2), decreases from 1.049 before the early opening to 1.026 afterwards for the two-year window around the early opening and from 1.062 to 1.029 for the one-year window

[53] Stoll and Whaley calculate the ratio of open-to-open to close-to-close return variance to be 1.13 for all NYSE stocks for the period 1982–1986.

(calculations based on the S&P 500). If the longer trading time reduces more of the opening noise (the price formation hypothesis is true), then the shorter closure time reduces equation (2) more than the result indicates. Thus, there is evidence that the extension of trading hours on the NYSE reduces transitory volatility in returns at the open relative to the close. Consistent with the findings of Gerety and Mulherin (1994), this evidence favors the price formation hypothesis and is against the trading mechanism type argument.

Assuming that the price formation hypothesis is true, a natural question that arises is whether the 2.7% decrease in non-trading time is sufficient to induce cleaner opening prices. There is reason to believe that it should since the 9:30 a.m. – 10:00 a.m. is a significant economic time. The cleaner opening prices probably reflect non-uniformity in noise creation during non-trading hours. We conjecture that noise is created primarily when information arrives and there is no opportunity to trade on it.[55] Between 9:30 a.m. - 10:00 a.m. most companies are open for business and, therefore, information arrival is likely to be high. Pre-early opening, the market is closed and there is no opportunity to trade on information arriving during 9:30 a.m. – 10:00 a.m. Thus, investors are unable to revise their interpretations of the information using reported prices. This could generate disproportionately more noise in prices when the market opened at 10:00 a.m.

---

[54] Stoll and Whaley (1990) find that closing prices are not totally void of noise.

2.7. CONCLUSION

We analyze the effect of extending trading hours on daily trading activity and daily return variability as well as on the observations that intraday volume and returns are U-shaped, trading time return variability is higher than non-trading time return variability, and transitory volatility is higher in opening prices than in closing prices. We find evidence consistent with an increase in trading activity when the NYSE extended its trading hours on September 30, 1985. In particular, the evidence suggests that increasing trading time can generate trade.

One possible explanation of this result is that the early opening eliminates any offsetting of overnight information that might occur between 9:30 a.m. and 10:00 a.m. Since most firms are open during this period for normal business, the rate of information arrival is likely to be high and, consequently, cancellations of accumulated overnight information are also likely to be high during this period. The predictions of this explanation are consistent with the significant increase in turnover observed after the early opening and, also, the insignificant increase observed when the closing period was extended from 3:30 p.m. to 4:00 p.m. on October 1, 1974. The information cancellation hypothesis suggests that the late closure in 1974 will not generate a significant increase in trading activity because accumulated information at 3:30 p.m., which follows continuous trading, is unlikely to be high. Thus, unlike the early opening, the late closure does not prevent significant information offsetting.

This explanation suggests that accumulated information at the beginning of the new trading period and the intensity of information arrival during this new trading period

---

[55] Hertzel et al. (1990) observe that in the foreign currency markets, there is concentration of noise when information arrival is high (that is, during the business hours of the country that the currency originates).

are crucial in determining whether a longer trading time will lead to extra trade. The result has significant implications for the move towards 24-hour trading. It suggests that trading activity may not increase substantially if firms are closed during the new hours and information arrival is low. The low after hours (4:00 p.m. to 6:30 p.m.) trading volume may be explained, at least in part, by the low level of accumulated information at 4:00 p.m. and the fact that most firms are closed after 5:00 p.m.

The study also finds that longer trading hours would not lead to an increase in daily return variability. Although we find a significant increase in daily return variability for the two-year window around the early opening, a number of factors suggest that this is *not* due to the longer trading hours. First, the increase in return variability is not consistent with an increase in private information production since trading time return variability does not increase relative to non-trading time return variability. This is the case even though a significant increase in trading activity is observed as early as one month after the early opening. Furthermore, there is no increase in daily variability for either the three-month window around the early opening or the late closure in 1974. This evidence and the findings of other researchers lead us to conclude that the significant increase observed over the two year window is due to other factors such as the instability that characterized the market prior to the October 1987 crash.

We also utilize this experiment to shed new light on some empirical facts in the microstructural literature and the hypotheses propose in the literature. In particular, we find that the intraday volume and return variability patterns become more U-shaped after the early opening. The data suggest that these results can be partially explained by shifts in trades from other periods to the open and the close. The observation that the intraday volume and return variability patterns become steeper at the open is inconsistent with the predictions of the Hong and Wang model. However, the increase in the convexity of intraday turnover and return variability is consistent with their predictions. The study also finds evidence consistent with the private information based explanation of why trading time return variability is higher than non-trading time return variability, which supports the findings of French and Roll (1986) and Barclay et al. (1990). Finally, the study finds that extending trading hours reduces transitory return volatility at the open relative to the close. This finding supports the price formation hypothesis and is consistent with the findings of Gerety and Mulherin (1994).

Table 2.1. Regression of Daily Turnover (with and without trend squared): Stocks: Two Years Around the Early Opening.

Variable	Estimate	Standard Error	Sign. Prob	Estimate	Standard Error	Sign. Prob
INTERCEP	0.418091	0.097242	0.001	0.399641	0.097731	0.001
TRN1	0.52006	0.032768	0.001	0.516105	0.032815	0.001
TRN2	0.025146	0.036715	0.493	0.022766	0.036704	0.535
TRN3	0.093335	0.036491	0.010	0.090969	0.03648	0.013
TRN4	0.093809	0.032471	0.004	0.090712	0.032488	0.005
DUMMY*	0.155817	0.056319	0.006	0.162909	0.056412	0.004
TUE	0.451498	0.038694	0.001	0.450873	0.038657	0.001
WED	0.379715	0.039928	0.001	0.379307	0.039888	0.001
THUR	0.351672	0.040282	0.001	0.351145	0.040242	0.001
FRI	0.264651	0.038765	0.001	0.264169	0.038727	0.001
FEB	0.009948	0.059[illegible]03	0.866	0.01279	0.059266	0.829
MAR	-0.03784	0.057535	0.510	-0.0369	0.05748	0.521
APR	-0.11593	0.058293	0.047	-0.1195	0.058271	0.041
MAY	-0.15723	0.059182	0.008	-0.16253	0.059202	0.006
JUN	-0.14608	0.059789	0.014	-0.15032	0.059779	0.012
JUL	-0.16403	0.059486	0.006	-0.16765	0.059463	0.005
AUG	-0.15222	0.059726	0.011	-0.15492	0.059686	0.010
SEPT	-0.16508	0.062753	0.009	-0.16746	0.062705	0.008
OCT	-0.21168	0.061873	0.001	-0.21879	0.061948	0.001
NOV	-0.17355	0.059886	0.004	-0.17641	0.059848	0.003
DEC	-0.07605	0.058576	0.194	-0.07873	0.058538	0.179
TREND	0.000275	9.76E-05	0.005	0.000568	0.000196	0.004
TREND SQ.				-2.8E-07	1.6E-07	0.085
R-square	0.7751			0.7758		
Adj. R-sq.	0.7703			0.7707		
F-Value	161[a]			154[a]		

The table reports the result of regressing turnover (TRN) on its lagged values (TRNi, where i=1,2,3,4 denote the i-th lag), DUMMY which is an indicator variable that takes a value of 0 pre-early opening and 1 after it, days of the week, months of the year, TREND (trend), and it square TREND SQ. Daily market turnover is calculated as the average of the turnover of the firms. The data is based on the sample of 2,676 firms from the CRSP database.

[a] significant at 1% level.

Table 2.2. Difference in Mean Turnover Tests: Three Months Before and Three Months After September 30th for the Years 1980 to 1990.

Year	Average Turnover (July 1 - Sept. 29)	Average Turnover (Sept. 30 -Dec. 31)	Percentage Increase	t-statistic	P-Value
1980	1.590	1.564	-2%	-0.471	.638
1981	1.220	1.254	2.7%	1.11	.267
1982	1.955	2.570	31%	5.08	.0001
1983	2.180	2.289	5%	1.757	.081
1984	2.369	2.241	-6%	-1.042	.299
1985	2.434	2.913	20%	4.489	.0001
1986	3.006	3.017	0.4%	0.108	.914
1987	3.254	3.557	9%	1.570	.119
1988	2.442	2.353	-4%	-1.173	.243
1989	2.572	2.517	-2%	-0.561	.576
1990	2.335	2.218	-5%	-1.25	.214

The table shows the average daily turnovers for periods July 1 – September 29 and September 30 – December 31 of 1980 through 1990. The t-statistics test the difference in these means for the different years. It can be observed that turnover in the three months after the early opening is significantly higher than the three months before. The control years display insignificant differences with the exception of 1982.

Table 2.3. Regressions of Daily Turnover of Stocks listed on the LSE and the NYSE and of Stocks listed on the NYSE but *not* the LSE.

	LSE and NYSE stocks		All other Stocks	
Variable	Estimates	Sign. Probability	Estimates	Sign. Probability
INTERCEP	1.019246	0.0001	0.411124	0.0001
TRN1	0.443802	0.0001	0.51924	0.0001
TRN2	0.010236	0.7688	0.027411	0.4558
TRN3	0.102906	0.003	0.092286	0.0117
TRN4	0.011554	0.7154	0.095523	0.0034
DUMMY	0.161913	0.0793	0.155065	0.006
TUE	0.517791	0.0001	0.450448	0.0001
WED	0.53096	0.0001	0.378769	0.0001
THUR	0.357444	0.0001	0.351632	0.0001
FRI	0.286566	0.0001	0.264681	0.0001
FEB	-0.17932	0.0785	0.01157	0.8452
MAR	-0.27949	0.0052	-0.03604	0.5309
APR	-0.25863	0.01	-0.11465	0.0493
MAY	-0.40883	0.0001	-0.15467	0.009
JUN	-0.41417	0.0001	-0.14356	0.0164
JUL	-0.41355	0.0001	-0.16177	0.0066
AUG	-0.44775	0.0001	-0.14938	0.0125
SEPT	-0.40624	0.0002	-0.1626	0.0096
OCT	-0.35541	0.0005	-0.20969	0.0007
NOV	-0.33131	0.0011	-0.17163	0.0042
DEC	-0.36782	0.0003	-0.07313	0.2117
TREND	0.000309	0.0589	0.000274	0.005
R-Square	0.4706		0.7764	
Adj. R-Sq.	0.4593		0.7716	
F-Value	41[a]		162[a]	

The table reports the results of the regression of turnover on our identified explanatory variables for stocks dually listed on the NYSE and LSE (London Stock Exchange) and that for those listed on the NYSE but not on the LSE. There were 65 stocks listed on the NYSE and LSE. The estimates of the coefficients on the DUMMY variable suggest that London did not lose trades to the NYSE after the inception of the early opening.

[a] significant at 1% level.

Table 2.4. Changes in the One-And-Half Hour turnover after the Early-Opening

Time	Change in Turnover
10:00 a.m. -11:30 a.m.	-2.4%
11:30 a.m.-1:00 p.m.	7.7%
1:00 p.m. - 2:30 p.m.	6.4%
2:30 p.m. – 4:00 p.m.	11.3%

To obtain the changes in turnover, we first calculate the mean turnover two years before the early opening and two years after it for each sub-period. We then calculate the percentage change in the mean turnover for each sub-period.

Table 2.5. Variance and Relative Variance of the One-and-Half Hour S&P 500 Returns

Time	Pre (1 yr)	Post (1 yr)	Rel. (pre)	Rel. (post)	t-statistics
10:00-11:30 a.m.	0.144	0.148	29%	18%	-2.33
11:30-1:00 p.m.	0.064	0.100	13%	12%	-0.212
1:00 – 2:30 p.m.	0.071	0.081	14%	9.8%	-0.891
2:30 – 4:00 p.m.	0.145	0.231	29%	28%	-0.212

All variances are multiplied by 10^4. The sub-period variances are relative to the close-to-close return variance. The t-statistics test the difference in the pre- and the post-early opening relative variances for each one-and-half hour of trade. For each year (for example, Sept. 30, 1994 – Sept. 29, 1985), we calculate the sub-period's return variance over a year and divide it by the daily return variance. This gives us four relative return variances for each year, and 16 relative return variances for the two-year window around the early opening. We then find the changes in relative variances over the years for each sub-period. These changes in relative variances are then used to calculate the standard deviation. The table reports the t-test statistics for the differences in relative return variances of Sept. 30, 1994 - Sept. 29, 1985 and Sept. 30, 1985 - Sept. 29, 1986 for each sub-period. The only significant decrease at 5% level of test occurs between 10:00-11:30 a.m. (the results of the other years are not reported).

Table 2.6. Variance Ratios to Estimate the Correlations among Intraday Returns

	1 Year Around EO		2 Years Around EO	
	Before	After	Before	After
$VR_{10:30}$	0.893	0.819	0.936	0.828
$VR_{11:00}$	0.703	0.910	0.728	0.878
$VR_{11:30}$	0.718	1.042	0.762	0.929
$VR_{12:00}$	0.883	1.149	0.852	0.968

The table reports the ratios of the variances of the close to period i return (where i is the subscript on VRi) divided by the sum of the variances of the close-open return and the 30-minute trading time returns between the close and period i [see equation (1)]. Thus, a variance ratio less than one indicates negative correlations among the returns in the interval and a ratio greater than one suggests positive correlations among the intraday returns.

Table 2.7. Variances and Relative Variances of Thirty-Minute S&P 500 Returns Two Years before the Early Opening and Two Years after it.

Time	Pre	Post	Prop (pre)	Prop (post)
9:30 – 10:00 am		0.2102		25.5%
10:00 – 10:30 am	0.1308	0.0523	26.3%	6.3%
10:30 – 11:00 am	0.0284	0.0378	5.7%	4.6%
11:00 – 11:30 am	0.0186	0.0396	3.7%	4.8%
Open-close	0.504	0.827	101.4%	100.2%
Close-close	0.497	0.825		
Open-open	0.523	0.848	104.6%	102.5%

All variances are multiplied by 10^4. Variances are relative to close-to-close return variance.

Table 2.8. Regression of Daily Turnover of all Stocks: Two Years Around the Late Closure

Variable	Estimate	Standard Error	Sign. Prob
INTERCEP	0.092728	0.039332	0.0186
TRN1	0.592164	0.031585	0.0001
TRN2	0.021987	0.036577	0.5479
TRN3	0.121719	0.036454	0.0009
TRN4	0.138119	0.031444	0.0001
DUMMY*	0.019811	0.026132	0.4486
TUE	0.088747	0.018562	0.0001
WED	0.054676	0.018819	0.0038
THUR	0.06303	0.018829	0.0008
FRI	0.01506	0.018768	0.4225
FEB	0.029781	0.029071	0.3059
MAR	-0.00416	0.028254	0.8829
APR	-0.00851	0.028999	0.7692
MAY	-0.00427	0.028965	0.8828
JUN	0.000463	0.029419	0.9874
JUL	-0.02166	0.02921	0.4586
AUG	-0.03365	0.03066	0.2727
SEPT	-0.0059	0.03064	0.8474
OCT	0.000349	0.028855	0.9903
NOV	0.013318	0.028906	0.6451
DEC	0.066102	0.028727	0.0216
TREND	-2.6E-05	4.63E-05	0.5678
R-square	0.7077		
Adj. R-sq.	0.7014		
F-Value	113[a]		

The table reports the results of the regression of turnover on the relevant explanatory variables after the inception of the late closure. Two years of data before and after the event were used. From the results, the coefficient estimate of the DUMMY variable is not significant.

[a]significant at 1% level.

Fig. 2.1. Thirty-Minute Turnover Two Years Before the Early Opening and Two Years After it

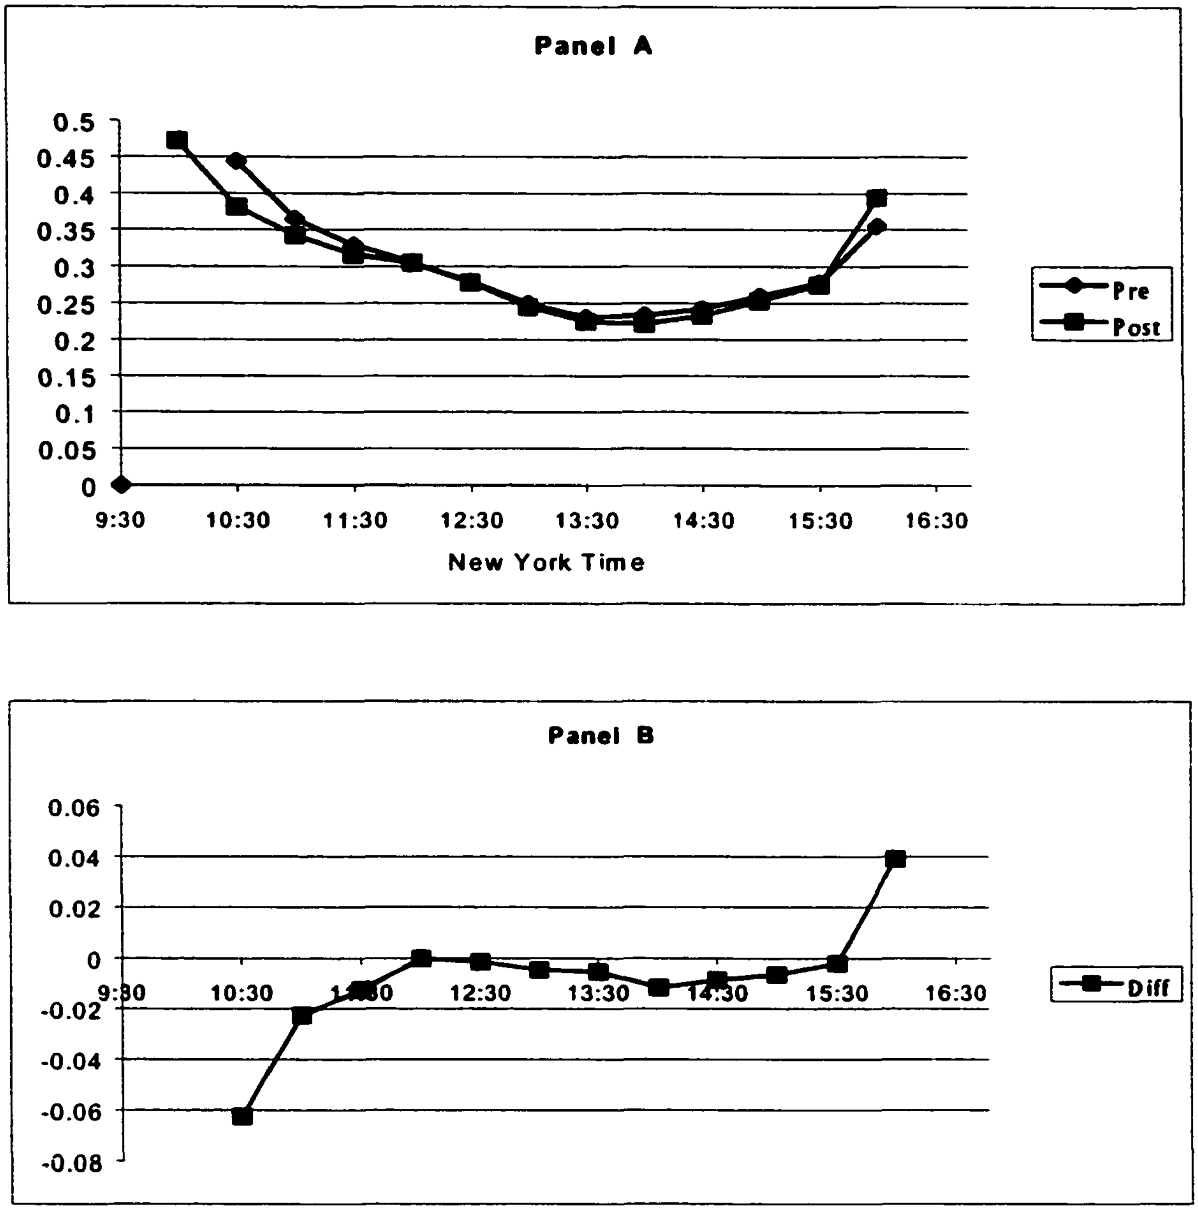

Panel A represents the mean of the 30-minute turnover pre- and post-early opening. The pre-early opening mean turnover for each sub-period is calculated as the mean of the turnover for the sub-periods from September 30, 1983 to September 29 1985 and the post-early opening period covers the sub-periods from September 30, 1985 to September 29, 1987. The post-early opening mean turnovers are scaled down such that the post and the pre-early opening mean turnovers are equal at 12:00 noon (turnover between 11:30 a.m. – 12:00 noon). The difference between the means is presented in Panel B.

Fig. 2.2. Thirty-Minute Turnover: Comparison of Trading time effects

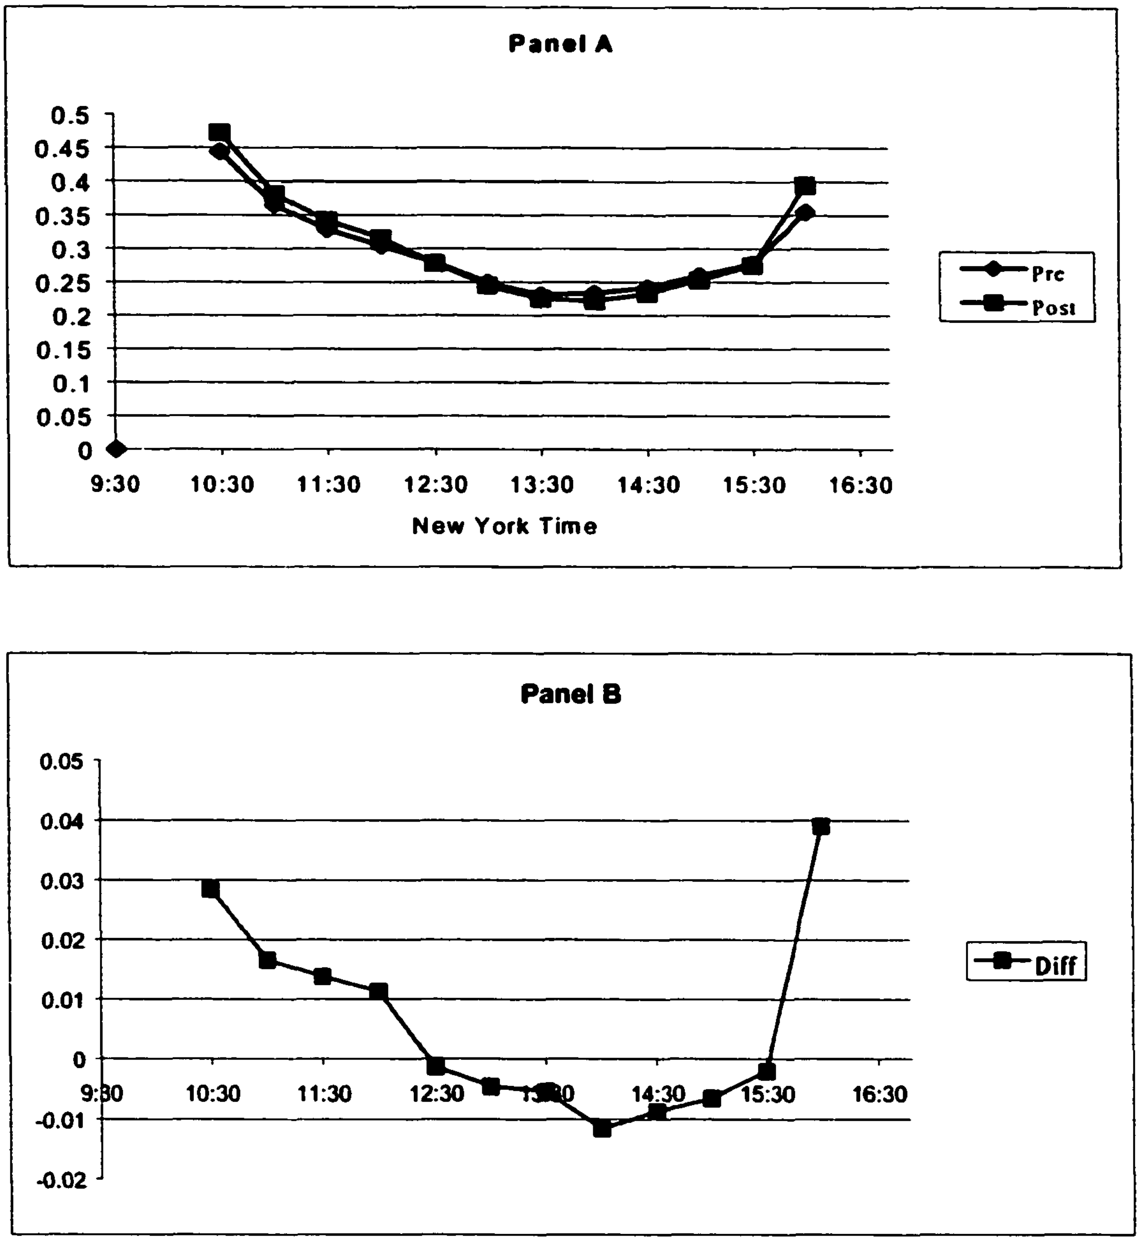

Panel A shows the pre- and the post-early opening turnover patterns. The turnover between 11:30 a.m. and 12 noon has been removed from the post-early opening so that trading periods can be compared. The difference between the turnovers is presented in Panel B. The convex shape of the difference curve near the open (two hours after the open) and near the close (two hours to the close) suggests that intraday turnover pattern becomes more U-shaped near the open and the close after the early opening.

Fig. 2.3. Thirty-Minute S&P 500 Return Variance

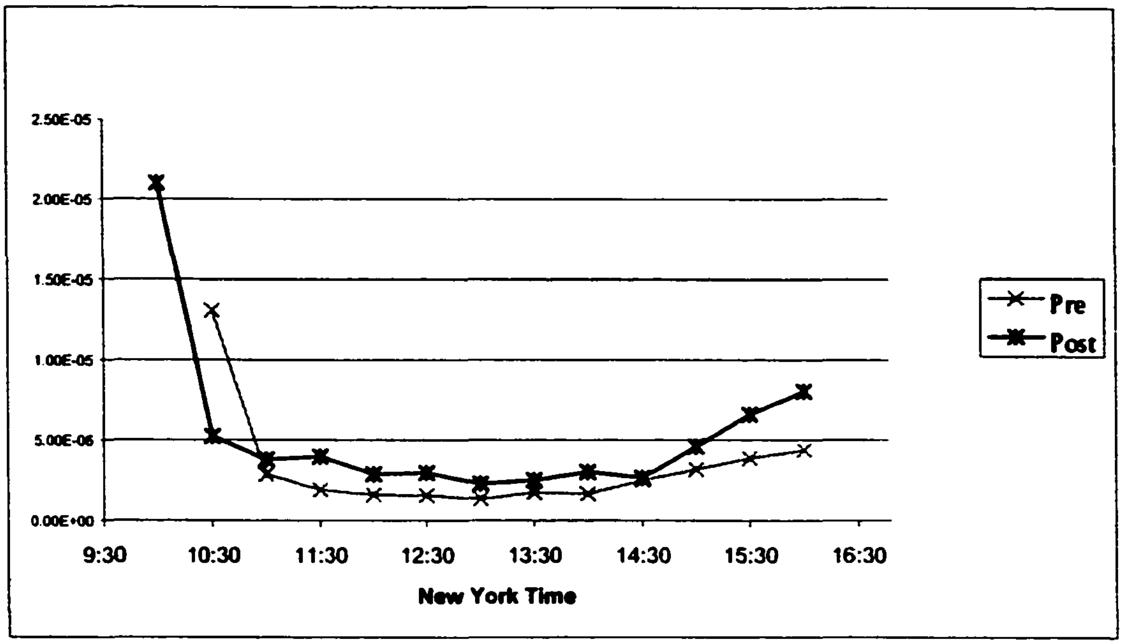

The figure shows the S&P 500 return variances for the overnight period and each thirty-minute of trade before and after the early opening. The pre-early opening return variances are calculated as the variance of the returns two years before the early opening for the appropriate intraday interval. The variances of the post-early opening returns are similarly obtained.

References:

Admati, A. R. and P. Pfleiderer (1988), A Theory of Intraday Patterns: Volume and Price Variability, *Review of Financial Studies*, 1, 3-40.

Amihud, Y., and H. Mendelson (1991) "Volatility, Efficiency and Trading: Evidence from the Japanese Stock Market," *Journal of Finance*, 46, 1765-1789.

______, (1987) "Trading Mechanisms and Stock Returns: An Empirical Investigation," *Journal of Finance*, 42, 533-553.

Barclay, M.J., Litzenberger, R.H., and Warner, J.B. (1990), Private Information, Trading Volume, and Stock Return Variances, *Review of Financial Studies*, 3(2), 233-253.

Black, F. (1986), Noise, *Journal of Finance*, 41, 529-543.

Berry, T. D., and K. M. Howe (1994), Public Information Arrival, *Journal of Finance*, 49(4), 1331-1346.

Booth, G., and M. Chowhury (1996), Information Noise and Stock Return Volatility: Evidence from Germany, *Applied Economics Letters*, 3, 537-540.

Chan, K. C., W. Fong, B. Kho, R. M. Stulz (1996), Information, Trading and Stock Returns: Lessons from Dually-listed Securities, *Journal of Banking & Finance*, 20, 1161-1187.

Chowhry, B. and V. Nanda (1991), Multi-Market Trading and Market Liquidity, *Review of Financial Studies*, 3, 483-511.

Dimson, E. (1979), Risk Measurement when Shares are Subject to Infrequent Trading, *Journal of Financial Economics*, 7, 197-226.

Dow, J and G. Gorton (1993), Trading, Communicating and the Response of Price to New Information, *Economic Journal*, 103, 639-646.

Fama, E. F. (1965), The Behavior of Stock Market Prices, *Journal of Business*, 38, 34-105.

French, K. and R. Roll (1986), Stock Return Variances: The Arrival of Information and the Reaction of Traders, *Journal of Financial Economics*, 17, 5-26.

Foster, D. and S. Viswanathan (1990), A Theory of the Interday Variations in Volume, Variance, and Trading Costs in Securities Markets, *Review of Financial Studies*, 3, 593-624.

Gerety, M. S. and J. H. Mulherin (1994), Price Formation on Stock Exchanges, The Evolution of Trading within the Day, *Review of Financial Studies*, 7, 609-629.

Granger, C. W. and O. Morgenstern (1970), *Predictability of Stock Market Prices*, Heath-Lexington, Mass.

Hertzel, M.G., C. S. Kendall and P. E. Kretzmer (1990), The Volatility of Asset Return during Trading and Nontrading Hours: Some Evidence from the Foreign Exchange Markets, *Journal of International Money and Finance*, 9, 335-343.

Hong, H. and J. Wang (2000), Trading and Return Under Periodic Market Closures, *Journal of Finance*, 55, 297-354.

Jain, P. and G. Joh (1988), The Dependence between Hourly Prices and Trading Volume, *Journal of Financial and Quantitative Analysis*, 23, 269-283.

Kleidon, A. and I. Werner (1996), U.K. and U.S. Trading of British Cross-listed Stocks: An Intraday Analysis of Market Integration, *Review of Financial Studies*, 9, 619-644.

Kyle, A. (1985), Continuous Auctions and Insider Trading, *Econometrica*, 53, 1315-1335.

Leach, C. and A. Madhavan (1993), Price Experimentation and Security Market Structure, *Review of Financial Studies*, 6, 375-404.

Oldfield, G. S. and R. J. Rogalski (1980), A Theory of Common Stock Returns over Trading and Non-trading Periods, *Journal of Finance*, 35, 729-751.

Romer, D. (1993), Rational Asset-Price Movements without News, *American Economic Review*, 83, 1112-1130.

Scholes, M. and J. Williams (1977), Estimating Betas from Nonsynchronous Data, *Journal of Financial Economics*, 5, 309-327.

Stoli, H. and R. Whaley (1990), Stock Market Structure and Volatility, *Review of Financial Studies*, 3, 37-71.

Tauchen, G. E. and M. Pitts (1983) "The Price Variability-Volume Relationship on Speculative Markets," *Econometrica*, 51(2), 485-505.

Wood, R. A., T. H. McInish, and J. K. Ord (1985), An Investigation of Transaction Data for NYSE Stocks, *Journal of Finance*, 40, 723-741.

Appendix

A model of Information Flow and the Volume of Trade

Following Tauchen and Pitts (1983), we assume that the market consists of J active traders, and within the day the market passes through a sequence of distinct Walrasian equilibria. At the time of the ith intraday equilibrium the desired position Q_{ij} of the jth trader is given by the linear function:

$$(1) \qquad Q_{ij} = \alpha[P_{ij}^{*} - P_i],\ \alpha > 0 \qquad\qquad (j=1,2,\ldots,J)$$

where P_{ij}^{*} is the jth trader's reservation price and P_i is the market price in period i. Given this, each Walrasian equilibrium will be characterized by $\sum_{j=1}^{J} Q_{ij} = 0$, and the market price, $P_i = \frac{1}{J}\sum_{j=1}^{J} P_{ij}^{*}$.

Suppose in period i, information arrives and changes the traders' reservation prices, then the resultant change in the market price and the associated volume will be given by equations (2) and (3).[56]

$$(2) \qquad \Delta P_i = \frac{1}{J}\sum_{j=1}^{J} \Delta P_{ij}^{*}, \quad \text{where } \Delta P_{ij}^{*} \equiv P_{ij}^{*} - P_{i-1,j}^{*}$$

$$(3) \qquad V_i = \frac{1}{2}\sum_{j=1}^{J} | Q_{i,j} - Q_{i-1,j} | = \frac{\alpha}{2}\sum_{j=1}^{J} | \Delta P_{ij}^{*} - \Delta P_i |$$

Assume that a day can be divided into i=1,...,c (where c denotes the last trading period). Furthermore, suppose the market is closed from periods 1 through 3 (for simplicity),[57] then daily volume will be given by V^B [equation (4)].

$$(4) \qquad V^B = \sum_{i=4}^{c} V_i = \frac{1}{2}\sum_{i=4}^{c}\sum_{j=1}^{J} | Q_{i,j} - Q_{i-1,j} | = \frac{\alpha}{2}\sum_{i=4}^{c}\sum_{j=1}^{J} | \Delta P_{ij}^{*} - \Delta P_i | ,$$

If the trading period is extended by opening the market in period 3, then daily volume after the extension will be given by V^A [equation (5)].

[56] For information to induce trade in this model, it is necessary that it generates different changes in reservation prices. Changes in traders' reservation prices can be different either because they receive different private information (heterogeneous information) or because they interpret public information (homogeneous information) differently.

[57] Also, as noted by Hong and Wang (2000), in terms of information arrival, the overnight period is a small fraction of the trading time period.

(5) $$V^A = \sum_{i=3}^{c} V_i = \frac{1}{2}\sum_{i=3}^{c}\sum_{j=1}^{J} | Q_{i,j} - Q_{i-1,j} | = \frac{\alpha}{2}\sum_{i=3}^{c}\sum_{j=1}^{J} | \Delta P_{ij}^{*} - \Delta P_i |$$

Our task is to find $V^A - V^B$. In order to do this, we need to make some specific assumptions about the nature of information flow. We assume the following:

(i) information flow is unchanged by the extension and accumulated information in the period *i-1* gets impounded in prices at period *i*.
(ii) only one piece of information can arrive at a given time interval and, in any interval, traders could get no information (N), good news (G) or bad news (B) about a particular stock.
(iii) for any interval *i*, information arrival (F) is independently identically distributed (iid) with the following probabilities: P(F=N) = 0.5, P(F=G) = P(F=B) = 0.25.
(iv) for any trader j, a piece of good news will cancel a piece of bad news (partial cancellations are not considered).

Let V_i^x for x = (A,B) denote volume in period i after and before the extension, respectively. Then, given assumption (i), the difference in volume after extending trading time is

$$V^A - V^B = E(V_3^A) + [E(V_4^A) - E(V_4^B)].$$

Before the extension

Although changes in the reservation prices will be influenced by information arriving a period before [assumption (i)], changes in the reservation prices at the open will depend on accumulated closing period and overnight information. Let F_j^k, denote investor j has accumulated k pieces of good or bad news at the open. We make this explicit by writing $\Delta P_j^*(F^k)$ and $\Delta P(F^k)$ as the changes in the reservation and market prices, respectively, in the opening period given that k=0,1,2,3,4 of good or bad news have accumulated. Thus, k is the net good or bad new that accumulates over the closing period. Given three periods of closure and one closing period, a maximum of four pieces of good or bad news can accumulate. Note that the changes in reservation prices at the open only depend on accumulated information and not the particular time that the exchange opens and, hence, we drop the intraday time subscript *i*.

Information arrival at each interval is iid with three possible outcomes [assumptions (ii) and (iii)] and, therefore, accumulated information at the open has a trinomial distribution, which depends on the length of the overnight period. The probability density function (pdf) for g pieces of good news and b pieces of bad news over the closing and the overnight periods is given by:

$$\begin{cases} \dfrac{4!}{g!b!(4-g-b)!}(.25)^g(.25)^b(.5)^{4-g-b}, & 0 \le g+b \le 4 \\ 0\ , & otherwise \end{cases}$$

Given this, we can calculate the probabilities of different levels of accumulated good or bad news at the open (k) by noting that a piece of good news cancels a bad one. For example, k=3 can be obtained by g=3 and no news in one period or b=3 and no news in one period. Since we are only interested in the volume effects, cases where g and b are equal are symmetric and are combined. Thus, the probabilities associated with the different levels of accumulated information are:

$P(k=4) = .0078125$
$P(k=3) = .0625$
$P(k=2) = .21875$
$P(k=1) = .4375$
$P(k=0) = .2734375$

The expected volume in the opening period is, therefore, given by:

$$(6) \qquad E\left(V_4^B\right) = \frac{\alpha}{2}\sum_{j=1}^{J}\left\{\begin{array}{l} .0078125*|\Delta P_j^*(F^4)-\Delta P(F^4)|+.0625*|\Delta P_j^*(F^3)-\Delta P(F^3)| \\ +.21875*|\Delta P_j^*(F^2)-\Delta P(F^2)|+.4375*|\Delta P_j^*(F^1)-\Delta P(F^1)| \end{array}\right\}$$

After an Extension:

After the an extension of trading time, the pdf of the trinomial distribution that describes information flow from the closing period to the open is given by:

$$\begin{cases} \dfrac{3!}{g!b!(3-g-b)!}(.25)^g(.25)^b(.5)^{3-g-b}, & for\ 0 \le g+b \le 3 \\ 0, & otherwise \end{cases}$$

and the different possible levels of accumulated information have the following probabilities:

$P(k=3) = .03125$
$P(k=2) = .1875$
$P(k=1) = .46875$
$P(k=0) = .3125$

$$(7) \qquad E\left(V_4^A\right) = \frac{\alpha}{2}\sum_{j=1}^{J} 0.5 * \mid \Delta P_j^*(F^1) - \Delta P(F^1) \mid$$

$$(8) \qquad E\left(V_3^A\right) = \frac{\alpha}{2}\sum_{j=1}^{J} \begin{Bmatrix} .03125 * \mid \Delta P_j^*(F^3) - \Delta P(F^3) \mid + .1875 * \mid \Delta P_j^*(F^2) - \Delta P(F^2) \mid \\ + .46875 * \mid \Delta P_j^*(F^1) - \Delta P(F^1) \mid \end{Bmatrix}$$

Difference in Volume:

$$(9) \qquad V^A - V^B = [\, E(V_3^A) \,] + [\, E(V_4^A) - E(V_4^B) \,]$$

$$= \frac{\alpha}{2}\sum_{j=1}^{J} \begin{Bmatrix} .53125 * \mid \Delta P_j^*(F^1) - \Delta P(F^1) \mid - .03125 * \mid \Delta P_j^*(F^2) - \Delta P(F^2) \mid \\ - .03125 * \mid \Delta P_j^*(F^3) - \Delta P(F^3) \mid - .0078125 * \mid \Delta P_j^*(F^4) - \Delta P(F^4) \mid \end{Bmatrix}$$

To simplify equation (9), we assume that absolute changes in reservation prices are linearly homogenous in the amount of accumulated information. That is $\mid \Delta P_j^*(F^k) \mid = k \mid \Delta P_j^* \mid$, where k is the amount of information accumulated and $\mid \Delta P_j^* \mid$ is the change in *j*'s reservation price based on a piece of information. This implies that

$$\mid \Delta P_j^*(F^k) - \Delta P(F^k) \mid = k \mid \Delta P_j^* - \Delta P \mid .$$

Given this, equation (9) simplifies to (10).

$$(10) \qquad V^A - V^B = \frac{\alpha}{2}\sum_{j=1}^{J} .34375 * \mid \Delta P_{1,j}^*(F^1) - \Delta P_1(F^1) \mid > 0$$

This completes the proof that volume increases after the extension. It should be noted that this result is robust to a number of simplifying assumptions which we used in the derivation. In particular, the positivity of equation (10) does not depend on:

(i) the linear homogeneity assumption used to simplify equation (9) and, indeed, if the changes in reservation prices are concave in the amount of information our results will be reinforced [equation (10) will be bigger positive value].

(ii) full cancellations, our result will hold to the extent that some good news can cancel bad ones and vice versa.

(iii) when information gets impounded in price, and the results will hold if current information is impounded in current prices

CHAPTER 3

KURTOSIS AND THE ACCURACY OF VALUE AT RISK

3.1. INTRODUCTION

Value at Risk (VaR) is the loss that will be exceeded over a pre-specified holding period, usually a day or two weeks, on some given fraction of occasions, typically 1% or 5%. The Bank for International Settlements (BIS), for instance, sets the confidence level at 99% over the next 10 days. The Derivative Policy Group proposes that over-the-counter derivative broker-dealers report the same VaR to the Securities and Exchange Commission (SEC). However, many firms use overnight VaR measure for internal purposes, as opposed to the two weeks standard commonly required by regulators.[1] Quantifying the VaR requires complete characterization of the future distribution of the risk factors. For example, if an institution has a position in 200 million German Marks, and is concerned about the exchange rate risk, then it will have to characterise the distribution of the changes in the exchange rate (German Mark vis-à-vis the U.S. dollar, for example). To do this, it may be assumed that changes in the exchange rate are normally distributed with a zero mean, and a variance that can be estimated from the historical returns. If past returns suggest that the daily standard deviation is 1%, for instance, then the daily VaR at 95% is 3.3m (1.65x1%x200m). In other words, the institution has 1-in-20 chance that the actual loss on the position will be greater than 3.3m German Marks in a day.

Various attempts have been made by researchers to obtain reliable VaR forecasts, which is essential for sound financial management. To do this requires forecasting the future distribution of the risks of the asset(s) accurately. This task has posed many challenges to researchers. The main distributional problem is that many financial returns, especially exchange rates, are leptokurtic relative to the Gaussian benchmark. That is extreme movements are more likely than the normal distribution predicts. The result of this is that observed daily portfolio losses have exceeded risk managers expectations.[2] It is, therefore, not surprising that researchers have made tremendous efforts to capture this salient distributional feature by employing different time series techniques. These include using stationary fat-tailed distributions such as the stable Paretian [McFarland et al. (1982)], Student's t [Calderon-Rossell and Ben-Horim (1982)], jump diffusions [Akgiray and Booth (1988)], and combining time-varying variances and fat-tailed independent shocks such as the Student's t/GARCH model of Baillie and Bollerslev (1989) and the jump diffusion/ARCH model of Jorion (1988). These methods have removed some (but not all) of the kurtosis in many financial returns. Recently, Barone et al. (1999) propose constructing empirical distributions within the GARCH framework to minimize the effect of imposing an incorrect distribution on the future risk structure.[3] This procedure, by construction, accounts for all the empirical kurtosis in the sample period.

Although research on the techniques to remove excess kurtosis is extensive, attempts have not been made to study the importance of the sources of kurtosis in accurately predicting VaR. In a GARCH setup, excess kurtosis can arise either from time-varying volatility or from fat-tail distributions. Attempts to model conditional kurtosis (kurtosis in errors from the mean equation) by the use of a fat-tail distribution changes the sharing of the modeled kurtosis between time-varying conditional variances

---

[1] For example, J. P. Morgan discloses its daily VaR at the 95% and Bankers Trust discloses its at 99%.
[2] Neftci (2000) and Longin (2000), for example, have proposed the use of extreme value theory to overcome this problem.

and fat-tail distributions and also increases the degree of the empirical kurtosis accounted for by the model.[4] Thus, the superiority of one method over the other could be due to the different degrees of the empirical kurtosis they capture or to the division of the explained kurtosis between time-varying variance and the assumed distribution. The technique of Barone et al., which combines parametric and non-parametric methods accounts for all the empirical kurtosis irrespective of the distributional assumption underlying the parameter estimates.[5] Thus, the approach eliminates differences in VaR forecasts that may be due to different degrees of modeling the unconditional empirical kurtosis. It, therefore, provides a framework within which we can study the relevance of the sources of the empirical kurtosis in forecasting VaR.

The Barone et al. methodology involves using scaled residuals from the mean equation to construct the future distribution of risk. For example, the Student's t or the normal distribution can be used to estimate the parameters of the GARCH equations and the resultant errors used to construct the future distribution of risk. If the unconditional kurtosis is largely driven by time varying volatility, then we would expect the normal distribution approach to yield superior VaR estimates. This is because, *ceteris paribus*, the normal distribution method would assign more time variation to the conditional variances since it has less density at the tails to absorb the empirical kurtosis than the t-distribution. That is, the normal distribution would explain more of the sample kurtosis via time varying variances than the t-distribution. More accurate normal distribution

---

[3] Various empirical distributions have been proposed for VaR estimations [for example, El-Jahel et al. (1999) constructed empirical distributions via matching moments], but they have not been used in a GARCH framework.

[4] For example, the use of a GARCH/t-distribution rather than a GARCH/Normal model will change both the estimated degree of heteroskedasticity of return (different variance equation estimates) and the amount of the empirical kurtosis explained by the model.

forecasts would, therefore, indicate that the unconditional kurtosis is mainly due to time variation in the variance rather than to fat-tail distribution *per se*. This would suggest that research for improved VaR forecasting methods should center on techniques that capture changing variances such as the use of more general stochastic volatility processes. On the other hand, more accurate t-distribution forecasts would suggest the empirical kurtosis is largely due to fat-tail distribution. In this case, the search for appropriate modeling techniques should focus on using fat-tail distribution techniques such as the use of mixtures of normals (jumps) and generalized error distributions.[6] Investigating the sources of the empirical kurtosis is essential because it lays the foundation and shapes the direction of future research on modeling the empirical kurtosis of asset returns. The study shows that, for the British Pound exchange rate, the unconditional kurtosis is mainly due to fat-tailed distribution, but for the other exchange rates, the source of the kurtosis appears to be unimportant.

The second task undertaken in this study is to study the relationship between modeling kurtosis and the accuracy of VaR estimates.[7] We use the normal/GARCH, the t-distribution/GARCH, and the Barone et al. approach to explore the relationship between successfully modeling kurtosis and the accuracy of VaR estimates. Among these techniques, the Barone et al. approach is most successful in modeling the empirical kurtosis, followed by the t-distribution, and the normal distribution is the least successful. Thus, by comparing the forecast performance of these modeling methodologies, we can extract the relationship between modeling the empirical kurtosis and the predictive power

---

[5] Indeed, Christoffersen (1998), in his study of evaluating forecasts, suggests that combining a dynamic variance specification with non-parametric error distribution will likely improve GARCH forecasts.
[6] Bate (1996) finds that jumps, rather than stochastic volatility, are able to explain the "volatility smile" evidence of implicit kurtosis of options on the deutsche mark.

of the resultant forecasts. This is important because it sheds light on the benefits of successfully modeling kurtosis in risk management. In particular, it provides lessons on whether the observed losses would be consistent with risk managers' expected losses if the empirical kurtosis is sufficiently modeled. We find that, as far as forcasting of VaRs of direct exposures are concern, modeling kurtosis as traditionally measured is inappropriate. For example, we find that all the exchange rates display excess kurtosis relative to the normal benchmark but the normal distribution approach yields the best forecast at 95% confidence level among the theoretical distribution forecasts across the currencies. As suggested by Duffie and Pan (1997), we find that the appropriate measure of kurtosis should be the number of standard deviations associated with the confidence level.

We conduct the study by using GARCH to model the returns on five major currencies and calculate the VaRs of the returns on the currencies. The exchange rates used are the Canada Dollar, Japanese Yen, British Pound, French Franc, and German Mark vis-à-vis the U.S. dollar. In particular, we estimate the models using five years of data (1990-1994) and assess their VaR forecasts using data from 1995 to May 2001.

The rest of the chapter is organized as follows. Section 3.2 discusses the motivation and background material. Section 3.3 presents the different evaluation criteria, and data is presented in section 3.4. Empirical analyses of the different techniques and their improvements in VaR calculations are explored in section 3.5. Section 3.6 summarizes the conclusions of the study.

---

[7] Most studies examine the reductions in kurtosis achieved by the various techniques without studying the potential improvements in VaR calculations [see, for example, Heish (1988), Baillie and Bollerslev (1989)]

3.2. MOTIVATION AND BACKGROUND MATERIAL

3.2.1. Review of Value at Risk

Risk management involves identifying, measuring, controlling, and communicating the risks taken by an institution to its senior managers. This study focuses on the searches to find accurate measures of risk. An accurate measure of risk is essential because it determines risk control and the level of risk that is communicated to senior management. Inaccurate risk measure would lead to improper risk control which would undermine financial management. Both underestimation and overestimation of risk are undesirable in risk management. For example, underestimation of risk would result in insufficient controls, which could result in unexpected financial losses. Overestimation of risk, on the other hand, would lead to unnecessary controls and result in losses of potential profits. Thus, the importance of the quest for accurate risk measurements cannot be overemphasized.

VaR has become a popular risk measurement tool for many financial institutions replacing the traditional sensitivity measures. This is mainly due to the introduction of RiskMetrics and some attractive features of VaR. First, VaR is a risk measure that can be applied to all traded products.[8] Therefore, it is a standard benchmark which allows risk being taken by different trading areas to be compared directly. Second, as VaR can be used to measure the risk on any product it can be combined across different trading areas to give a single figure for the risk being taken by all trading areas combined. This single number has become very important in view of the proposals for disclosure of financial risk which calls for firm wide risk measures. During the past decade, discussions

---

[8] The RiskMetrics is a linear VaR model based on variance-covariance of past security returns introduced by J.P. Morgan in 1993.

concerning risk disclosures by many organizations have made it clear that risk disclosures are important in financial reporting.[9] This consensus is largely driven by the significant losses reported due to increased usage of derivative instruments in the first part of the last decade.[10] The regulatory bodies responded with both conceptual guidelines and concrete recommendation and requirements. In January 1997, the Securities Exchange Commission (SEC) issued Financial Reporting Release No. 48 (FRR 48), which required (not just recommend), disclosures of risk related measures. One of the SEC recommended methods of risk disclosure is VaR.[11]

Third, the probability of a loss exceeding VaR is known and VaR also takes account of the correlations among asset returns.[12] Knowing the probability of a loss exceeding the VaR helps management to determine whether the level of risk is tolerable, and capturing the correlations among asset returns is crucial in accurately estimating and forecasting VaR of a portfolio. For a portfolio of many assets, it is a difficult statistical task to capture the correlations.[13] The advantage of VaR, however, is that it accommodates the correlations among assets. These features, together with the introduction of Riskmetrics in 1993, are the cornerstone of the wide acceptance of VaR as a risk measurement tool. Indeed, it is now a global *de facto* risk measurement standard and is expected, if not required, by most regulatory bodies in the G10 countries to be used

---

[9] Among others, these organizations include the Association for Investment Management and Research (AIMR), the American Institute of Certified Public Accountant (AICPA) and the Financial Executive Institute (FEI).

[10] For example, Barings went bankrupt in 1995 due to $1.3 billion loss from derivatives trading. Other examples include Metallgesellschaft ($1.3 billion) in 1993 and Orange County ($1.64 billion) in 1994.

[11] SEC required disclosure formats are: 1. Tables presenting the fair value of the instruments and sufficient information to determine the cash flow amounts expected by maturity dates. 2. Sensitivity analyses describing the effect on earnings, cash flows or fair value of changes in market rates or prices. 3. VaR disclosures on earnings, cash flows or fair value of the instruments from changes in market risk factors.

[12] The traditional sensitivity-based measures of risk cannot achieve these.

[13] For many assets, the correlation matrix becomes big and may not be invertible (singularity problems) and the estimates may be biased when the correlations among assets are non-stationary.

to report risk levels. In addition, some regulatory bodies now allow banks to use VaR as the basis for calculating the banks' regulatory capital requirements.[14] Also, rating agencies expect banks to have implemented comprehensive VaR systems and to use VaR measures to adjust their performances.

VaR, however, is not a panacea. It cannot, for example, effectively measure the market risk when the market is not behaving normally. VaR, therefore, seems to be ineffective when needed most. Another weakness of VaR is that it does not tell us how big the loss that exceeds the VaR will be. Thus, comprehensive risk measurement systems use stress testing to complement VaR measures in order to obtain more accurate measures of risk during market anomalies and to assess the maximum possible losses.

The widespread use and acceptance of VaR as a risk measurement and control tool from the middle of the 1990 calls for more research on VaR methodologies and finding ways that would make the estimates from VaR more efficient (see Glasserman (2000) on the quest for precision in risk estimates through the use of VaR). The contribution of this study is in line with this research agenda. In particular, we examine the role of the source of the kurtosis in accurately estimating VaR and also the precision gained in VaR estimates by modeling the empirical kurtosis.

3.2.2. Some Empirical Measurement Issues

Some critical empirical issues in calculating VaR are what distributional hypotheses are consistent with the observed properties of the assets' returns. The evidence from both time series analyses and options on currencies suggest that foreign exchange return volatilities cluster over time and their distributions are leptokurtic [see, for

[14] See Bank for International Settlements (1996).

example, Heish (1989)].[15] The extensive application of ARCH and GARCH techniques to model exchange rate returns reflects their ability to capture these empirical regularities.[16] Changing implicit conditional volatilities of exchange rates over time supports volatility clustering, and Bates (1994) finds excess kurtosis in a model-specific implicit distribution of options on the deutschemark (DM) and yen futures. The evidence regarding skewness is more mixed with time series estimates sensitive to the currency and the period used. For example, Bates (1996) finds that the DM exhibits substantial positive skewness during 1984-1985 and a non-stable skewness thereafter.

While there may be various theoretical reasons for "fat-tails" in empirical distribution of assets' returns, Duffie and Pan (1997) identify jumps and time-varying volatility as the probable causes. GARCH models can accommodate both of these sources of tail-fatness in the empirical distributions. Duffie and Pan, for example, argue that the impact of jumps can be replicated by mixing normals with different variances.[17] Thus, drawings from fat-tail distributions are consistent with jump diffusion models over the short term, and hence the use of an appropriate fat-tail distribution can capture the impact of jumps.[18] GARCH models are, therefore, natural procedures for investigating the sources of fat-tailness as the methodology accommodates both fat-tail distributions and time-varying variances.

In estimating VaR, three broad GARCH methodologies can be identified: the traditional GARCH which assumes errors from the mean equations are Gaussian,

---

[15] Heish (1988), for example, rejected the test of equal monthly variance for all five major currencies he considered from 1974 to 1983.

[16] See Bollerslev, Chou, and Kroner (1992) for a survey of the ARCH/GARCH literature, including their applications to foreign exchange rates.

[17] Apart from this, jump terms can be directly introduced in the GARCH equations (see, for example, Jorion (1988) and Bates (1996) where jump terms are included in the mean equations)

GARCH methods which assume the errors have conditional fat-tail distributions [for instance, the GARCH/Student's-t model of Baillie and Bollerslev (1989)], and the combined empirical distribution and time-varying variance of Barone et al. (1999). The mean and the variance equations of a GARCH model can be written as equations (1) and (2), respectively.

$$r_t = f(X) + \varepsilon_t \quad \text{and} \quad \varepsilon_t / \Omega_{t-1} \sim D(0, h_t) \tag{1}$$

$$h_t = a_0 + \sum_{i=1}^{p} \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^{q} \beta_j h_{t-j} \tag{2}$$

where r_t is the return at time t, f is some function, and X is a set of explanatory variables. The error from the mean equation (ε_t) has distribution D with mean 0 and variance h_t conditional on information up to time t-1 (Ω_{t-1}). Following Bollerslev (1986), the conditional variance is given by the GARCH model of orders p and q [GARCH(p,q)].

The GARCH/normal approach assumes that ε_t is normal given h_t and, consequently, the future distribution of risk is normal conditional on the future volatility. Given the conditional normality assumption, the parameters of the model are estimated and the variance equation is used to forecast the future variance of the asset return. Conditional on the estimated future variance, the distribution of future risk is normal and this allows the calculation of VaR at some level of confidence. Empirical results using the conditional normality approach, however, show that the errors of exchange rates returns are fat-tailed conditional on the variance [see Heish (1989) and Baillie and Bollerslev (1989)]. These results suggest that, given the assumed variance structure, the conditional normality assumption is inappropriate. The implication of this is that the modeled

[18] Bates (1996), for example, finds that stochastic volatility is unable to explain the volatility smile of the DM evidence of implicit excess kurtosis, but finds that jump fears can explain the smile.

kurtosis captured by the dynamics of the variance (equation 2) and the theoretical kurtosis of the normal distribution is insufficient to account for the unconditional sample kurtosis.

To tackle the conditional kurtosis problem evidenced in the Gaussian assumption, many researchers use different fat-tailed distributions to estimate the GARCH parameters. For example, Heish (1989) experimented with various non-normal distributions (t-distribution,[19] normal-Poisson mixture,[20] normal-lognormal distribution, generalized error distribution), and finds that exponential GARCH and any of the four non-normal distributions fits Canadian Dollar, the German Mark, and the Swiss Franc quite well. A normal-Poisson mixture fits the Japanese Yen, but these fat-tail distributions could not explain the excess kurtosis in the British Pound. Also, Baillie and Bollerslev (1989) find that the GARCH/t-distribution could not account for the excess kurtosis in the French Franc. The finding of conditional excess kurtosis is consistent with the U-shaped pattern of implicit volatilities across different strike prices, or the "volatility smile", which suggests conditional leptokurtosis in the underlying prices.[21]

To address the conditional kurtosis problem in using both normal and assorted fat-tail distributions, and also avoid imposing a particular distributional assumption on the structure of future risk, Barone et al. propose constructing empirical distributions via simulations. Their method involves drawing randomly from standardized mean errors ($\frac{\varepsilon_i}{\sqrt{h_i}}$, for i = 1,...,t), which are independent identically distributed (iid) random variables

---

[19] The t-distribution is a normal-inverted gamma mixture.
[20] This is a mixture of standard normal and a Poisson jump process, which yields the standard jump models.
[21] Cao (1992) and Shastri and Wethyavivorn (1987) document the volatility smile present in foreign currency option prices.

under the GARCH hypothesis.[22] These random samples are multiplied by the forecast standard deviation ($\sqrt{h_{t+1}}$) to adapt them to current volatility conditions. The scaled-standard errors are used to retrieve the forecast prices of the asset via the mean equation. This procedure is replicated a large number of times, for instance 5,000, and the resulting prices are used to construct the empirical distribution for calculating the VaR.

By using the standardized residuals from the mean equation to predict the future distribution of the asset, the methodology accounts for all the empirical kurtosis in the sample period. The approach is, therefore, suitable for studying the relationship between modeling excess kurtosis and the accuracy of VaR estimates. In addition, the method can be used to study the importance of the source of kurtosis in predicting the VaR. This is because the procedure accounts for all the empirical kurtosis of the sample regardless of the distributional assumption underlying the parameter estimates. Different distributional assumptions would, therefore, only alter the proportion of kurtosis explained by time-varying variance and the forecast empirical distribution. In this study, we estimate the parameters using the normal and the t-distribution (a representative of fat-tail distribution). Dynamic variance forecasts from normal likelihood and t-distribution likelihood would, in general, be different. This will generate differences in the kurtosis attributable to changing variances. Comparison of VaR forecasts from a normal likelihood *cum* forecast empirical distribution and a t-distribution likelihood *cum* forecast empirical distribution would shed light on whether the empirical kurtosis is more appropriately modeled by a dynamic variance or a fat-tail distribution.

[22] Note that iid assumption may be violated to various degrees empirically.

3.3. EVALUATION PROCEDURES

3.3.1. Operational Evaluation

There are different ways of evaluating VaR forecasts. Statistical evaluation procedures such as the root mean square error (RMSE) and likelihoods assess the ability of the model to forecast the center of the return distributions, but it is the accurate predictions of the tails that are important in VaR forecasts. A volatility-forecasting model will have a high likelihood/low RMSE if most of the returns on the test set lie in the center range of the predicted distribution. But since VaR models attempt to predict the worst-case scenarios, it is really the lower percentiles of the predicted distributions that should be examined.[23] In view of this, we use a variety of methods which have been proposed and used by various researchers [see, for example, Davé and Stahl (1997) and Lopez (1996)], and supplement these with statistical evaluations.

3.3.2. Observed and Predicted Exceedence

For any financial series $\{y_t\}_{t=1}^{T}$, information up to period t = s-1 can be used to estimate a model. Based on the estimated model, VaR forecasts at a given confidence p $[L_{t/t-1}(p),\infty]_{t=s}^{T}$ for long positions and $[-\infty,U_{t/t-1}(p)]_{t=s}^{T}$ for short positions can be calculated for the remaining periods (s to T).[24] The number of times the actual losses exceed the forecast VaR for both short and long positions can be obtained by defining an indicator function I_t as follows:

$$I_t = \begin{cases} 0, & if\ \ y_t \in [-\infty,U_{t/t-1}(p)] \\ 1, & if\ \ y_t \notin [-\infty,U_{t/t-1}(p)] \end{cases}, \qquad \text{for short positions, and}$$

[23] This in part explains why some researchers have proposed the use of extreme value theory.

$$I_t = \begin{cases} 0, & if\ y_t \in [L_{t/t-1}(p), \infty] \\ 1, & if\ y_t \notin [L_{t/t-1}(p), \infty] \end{cases}, \qquad \text{for long positions.}$$

The sum of the indicator variable, $\sum_{t=s}^{T} I_t$, therefore gives the number of times the actual losses exceed the VaR forecasts (observed exceedence) for a given coverage p. Many evaluation procedures are based on this sum. The Bank for International Settlements (BIS), for example, based its zone categorization (Green, Yellow and Red) of models on it. The Bank recommends that, on any day t, a model is in the Green zone if the sum of the indicator function over the last 250 days is not greater than four at 99% coverage. A model is in the Yellow zone if the sum of the indicator function is between five and eight (inclusive), and it is in the Red zone if it exceeds eight. Another evaluation method based on the sum of the indicator function is the Observed/Predicted exceedence ratio [see, for example, Davé and Stahl (1997)]. This performance measure is obtained by dividing the observed exceedence count by the expected exceedence. For a given coverage p, the expected exceedence is given by (1-p)(T-s). The Observed/Predicted exceedence ratio is thus given by $\sum_{t=s}^{T} I_t$ /(1-p)(T-s). An exceedence ratio greater that one indicates than the model underestimates risk, and a value less than one suggests that the model over-predicts risk.

The Observed/Expected exceedence ratio, like the BIS recommendation, is not sensitive to the degree of exceedence and hence cannot differentiate between models with the same exceedence count but different degrees of prediction errors. Davé and Stahl use the mean log likelihood, which is designed to measure the degree to which the losses

---

[24] $[-\infty, U_{t/t-1}(p)]_{t=s}^{T}$ and $[L_{t/t-1}(p), \infty]_{t=s}^{T}$ are respectively the forecasts of the upper and lower bounds of confidence intervals with significant level 2p based on information up to t-1.

exceed those predicted at a certain confidence level. The measure is given by the mean of the log-likelihoods of all events for which the observed loss exceeds the VaR. Thus, if Li (for i=1,...,k) denote the log-likelihoods of events for which the observe losses exceed the VaR predictions, then the measure is given by $\sum_{i=1}^{k} Li/k$. The degree of exceedence can also be measured by adding up all the losses that exceed the VaR forecasts. In this case, if Zi is the loss that exceeds the VaR for day i, then this measure will be given by $\sum_{i=s}^{T} Z_i$. However, a model with a smaller degree of exceedence may not be superior. This is because, theoretically, some violations are expected and hence some degree of exceedence must also be expected. Thus, observed degree of exceedence must be compared to some theoretical benchmark to ascertain a superior modeling technique.

A weakness in the evaluation procedures outline above is that they test for unconditional forecasts. Christoffersen (1998) argues forcefully that this is inappropriate in setups where non-trivial dynamics exists in the higher order moments of a series. Thus, in GARCH frameworks, where dynamic relationships exist among the variances, conditional evaluation procedures will be more useful in assessing the performances of the various models. One of the motivations of the Engle (1982) research is to predict dynamic intervals around point predictions. The insight was that the intervals should narrow in tranquil times and widen in volatile periods so that the occurrences of observations outside the interval forecast would be spread out over the sample and not in some clustered fashion. Davé and Stahl recognize this and, in an *ad hoc* fashion, use serial exceedence ratio to test a model's ability to capture outliers that may cluster. In particular, they use serial exceedence ratio to test a model's propensity for consecutive prediction failures. Serial exceedence count is obtained by counting the number of consecutive prediction failures. An indicator function for counting consecutive prediction failures S_t can be defined as follows:

$$S_t = \begin{cases} 1, & if\ \ I_t = 1\ \ and\ \ I_{t-1} = 1 \\ 0, & otherwise \end{cases}$$

The expected number of consecutive failures at coverage p is given by $(1\text{-p})^2(\text{T-s})$, and hence the serial exceedence ratio is $\sum_{t=s}^{T} S_t /(1\text{-p})^2(\text{T-s})$. If this ratio is greater than 1, then the model under-predicts the degree of dynamics in the higher order moments of the series. A ratio less than one suggests that the model over-predicts the degree of dynamics.

Christoffersen (1998) proposes a test for evaluating conditional interval forecasts (Lopez (1996) applies this test), which is crucial when higher order moment dynamics are present as suggested by GARCH models. An interval forecast that fails to take account of higher order moment dynamics may be correct on average in predicting the unconditional coverage, but in any given period it will have incorrect coverage characterized by clustered outliers. Thus, the previous performance measures are still valid if they are used to evaluate forecasts over a long horizon.

Interval forecasts are efficient with respect to information set Ω_{t-1}, if the $E[I_t/\Omega_{t-1}] = 1\text{-p}$ for t = s,...,T. The traditional method of assessing interval forecast [see, for example, Ballie and Bollerslev (1992) and McNees (1995)] is to assume that the information set is null and test $E[I_t] = 1\text{-p}$ by comparing $\hat{p} = 1 - \sum_{t=s}^{T} I_t /(T-s)$ to the true coverage p. But the presence of higher order moments dynamics such as time-varying variances suggest that testing the conditional accuracy of the forecasts is important. To demonstrate this, suppose information set Ω_{t-1} contains only values of the past indicator function, then $E[I_t/\Omega_{t-1}] = E[I_t/ I_{t-1}, I_{t-2}, I_{t-3}, \ldots, I_s] = 1\text{-}p$, for $t = s,\ldots,T$. The standard way of evaluating interval forecasts is equivalent to testing whether the indicator variable is an independent identically distributed Bernoulli random variable with probability p of being equal to zero and (1-p) of being equal to one. This test can be performed using the likelihood ratio test. The likelihood ratio statistic $LR_1 = -2\log[L(p; I_s, I_{s+1},\ldots, I_T)/L(\hat{p}; I_s, I_{s+1},\ldots, I_T)]$ has an asymptotic $\chi^2(1)$ distribution.

It can be seen that this statistic does not have the power to test whether the values of the indicator variable cluster in some dynamic fashion. In order to do this, Christoffersen assumes that the indicator function I_t follows a first-order Markov process. In this situation, the probability of the indicator variable is measurable with respect to its immediate past value. Accordingly, the likelihood function of the transition probabilities $\pi_{ij} = P(I_t = j/I_{t-1} = i)$, where $(i,j) = 0,1$ is given by

$$L_1 = L(\Pi_1; I_s, I_{s+1},\ldots, I_T) = (1-\pi_{01})^{n_{00}} \pi_{01}^{n_{01}} (1-\pi_{11})^{n_{10}} \pi_{11}^{n_{11}},$$

where n_{ij} is the number of occurrences of $I_t = j$ given $I_{t-1} = i$. If the sequence of the values of the indicator function is independent (i.e. $\pi_{ij} = \pi_{jj}$), then the likelihood function is given by

$$L_2 = L(\Pi_2; I_s, I_{s+1},\ldots, I_T) = (1-\pi_2)^{(n_{00}+n_{10})} \pi_2^{(n_{01}+n_{11})}.$$

π_{ij} can be estimated by $\hat{\pi}_{ij} = \sum_{t=s}^{T} (I_t = j / I_{t-1} = i)/(T-s)$ and the likelihood ratio statistic $LR_2 = -2\log[L(\hat{\Pi}_1; I_s, I_{s+1},\ldots, I_T)/L(\hat{\Pi}_2; I_s, I_{s+1},\ldots, I_T)]$, which is asymptotically $\chi^2(1)$ can be used to test the independence of the sequence of the indicator variable I_t. The likelihood ratio LR_2 does not depend on the confidence level p. To test both the coverage p and the independence of the indicator series, the likelihood ratio statistic

$LR_3 = -2\log[L(p; I_s, I_{s+1}, \ldots, I_T) / L(\hat{\Pi}_1; I_s, I_{s+1}, \ldots, I_T)]$, which is $\chi^2(2)$ can be used.[25]

We use the BIS criteria, the exceedence ratios, the degree of exceedence, and the tests suggested by Christoffersen to evaluate the VaR predictions of the different models we consider in this research.

3.4. DATA AND EMPIRICAL ESTIMATIONS

The empirical work requires the spot exchange rates of the currencies the different countries vis-à-vis the U.S. dollar. Daily data consisting of exchange rates for five major currencies between January 1, 1990 and May 31, 2001 were extracted from Datastream International. The exchange rates are defined as the amount of foreign currency one U.S. dollar can buy, and they are the rates prevailing at mid-day Eastern Time. The currencies are the Canadian dollar (CD), the Pound Sterling (BP), German Mark (DM), French Frank (FF), and the Japanese Yen (JY). The total number of trading days covered by the data is 2979. The in-sample period consists of 1305 data points (January 1, 1990 to December 31, 1994), and 1660 data points are use as the out-of-sample data (January 1, 1995 to May 14, 2001). I drop the last fourteen data points in order to avoid factional expected violations at 5% level of tolerance.

[25] See Christoffersen (1998) for details on the derivation of the test statistics. The likelihood ratio statistic $LR_3 = LR_1 + LR_2$.

3.4.1. Summary Statistics of Exchange Rate Returns

Exchange rates, E_t, are transformed to continuously compounded percentage returns: $r_t = 100 \times \log[E_t/E_{t-1}]$. Summary statistics of the historical behaviour of these returns are reported for the sample period (1990-1994) in Table 3.1. The results confirm the excess kurtosis found in most other studies [for instance, Heish (1989) and Baillie and Bollerslev (1989)]. Kurtosis ranges between 7.9 for the JP to 4.9 for the DM, and the excess kurtosis is significant at 1% for all the exchange rates relative to the normal value. From the table, the evidence on skewness is rather mixed.[26] The BP, the DM and the CD are positively skewed, but the FF and the JY are negatively skewed. The skewness is significant at 5% for the BP, the CD, and the JY. The excess kurtosis across the currencies and the significant skewness in three of them suggest that the unconditional distributions of the exchange rate returns are not normal. This is confirmed by the chi-square and the Jacque-Bera Lagrange Multiplier tests, which show strong rejection of the hypotheses that the unconditional distributions of the exchange rate returns are normal.[27]

We also test the hypotheses that the daily log-difference exchange rates for each year come from the same distribution over the years by using non-parametric test. In this test, the sample is divided into groups by the years and the empirical densities are compared. The Kruskal-Wallis test statistics (chi-square approximation) are also reported in Table 3.1. The results reject the null hypotheses that the yearly distributions of returns are the same for each currency.

---

[26] This is consistent with many studies that show the exchange rates do no display any consistent skewness pattern [see, for example, Jorion (1985)].

[27] The Jarque-Bera test statistic (LM) is given by $LM = N\left[\frac{g_1^2}{6} + \frac{g_2^2}{24}\right]$, where g_1 and g_2 are coefficients of skewness and kurtosis respectively.

The out-of-sample statistics are shown in Table 3.2. The exchange rates display significant excess kurtosis and, indeed, the excess kurtosis is higher in forecast period than in the sample period for all the currencies. The JY registered the highest increase while the BP recorded the least increase. Thus, the data allow us to study the performance of forecasting VaR, when the tails of the future distributions are getting fatter. From this, we can infer the performance of the VaR models when the future distributions are becoming less leptokurtic. Skewness, on the other hand, decreased for all the exchange rates in the forecast period relative to the sample period. This resulted in negatively skewed distribution for all the exchange rates in the forecast period. We can, therefore, study the forecasting performance of VaRs as skewness decreases and can make inferences about performance as skewness increases. Like the sample period, the Jarque-Bera and the chi-square tests reject the hypotheses that the exchange rate returns are from normal distributions. Likewise, the Kruskal-Wallis tests reject the hypotheses that exchange rates come from the same yearly distributions.

Several studies have reported that exchange rate returns have little autocorrelation but are strongly heteroskedastic [see, for example, Melino and Turnbull (1990)]. Table 3.3 shows the estimates of the first 12 autocorrelations coefficients, $\rho_r(k)$ and $\rho_{rr}(k)$, and the modified Box-Pierce statistics, $Q_r(K)$ and $Q_{rr}(K)$, for r_t and r_t^2 respectively. Under the null hypothesis that returns are independent identically distributed (iid), McLeod and Li (1983) demonstrate that $\rho_{rr}(k)$ has an asymptotic standard normal distribution and that $Q_{rr}(K)$ has an asymptotic chi-square distribution with K degrees of freedom. The modified Box-Pierce statistics are all insignificant for the returns on the exchange rates [highlighted], except for the British pound. The same statistics are all significant for the squared returns [highlighted], with the exception of the Yen.[28] This suggests that there is substantially more autocorrelation in the squared returns of the exchange rates, which indicates non-trivial dynamics among the variances of the exchange rate returns. This evidence supports the use of GARCH modeling techniques to capture the relationships among the variances of the exchange rate returns.

It is useful at this point to state three stylized facts about exchange rate returns, which are confirmed here. First, the direction of daily exchange rate returns is virtually unpredictable. With the exception of the BP, which shows some significant autocorrelations (see the Box-Pierce statistic), the means and the autocorrelations of the exchange rate returns are all insignificant. This suggests that historical returns do not provide any useful information for predicting the exchange rate returns. Brooks (1996), for example, finds that only very modest improvements are achieved by forecasts of various linear and non-linear univariate time-series models over random walk forecasts for daily Sterling exchange rate. Second, squared exchange rate changes have nontrivial dynamic structure. This is illustrated by the significant autocorrelations among the squared returns and their significant modified Box-Pierce statistics for all the currencies with the exception of the JY. Third, the distributions of the daily exchange rates have fairly fat tails relative to the normal distribution. The significant excess kurtosis reported in tables 3 and 4 (values in excess of 3) confirm this. Accordingly, Melino and Turnbull (1990) suggest that exchange rate models should take account of these facts. Since GARCH models are capable of capturing these empirical regularities, they have been used extensively to model exchange rates and we use them here for our investigations.

[28] Christoffersen (1998) also finds that there is not much evidence of conditional variance dynamics in the Swiss franc for the period 1988–1995, but the DM, the BP and the JY display different variance dynamics

3.5. THE GARCH MODELS

In the GARCH estimations, we distinguish between the distributional assumption used to estimate the parameters and the one used to forecast the future risk distribution. This distinction is important because, in the bootstrapping approach, the non-parametric empirical distribution used to forecast the future risk structure will, in general, be different from the distributional assumption used to estimate the parameters. We estimate the GARCH parameters using the normal and the t-distribution likelihoods and, consistent with the GARCH hypothesis, we use these distributions to project the future risk structure. In addition, we use the empirical distributions of the adjusted standardized residuals to forecast the distribution of the exchange rate returns under the different theoretical likelihood assumptions.

3.5.1. The Normal Likelihood

3.5.1.1. Estimation

Although different orders of GARCH have been proposed, the GARCH(1,1) is found to be adequate for many financial series including exchange rates [see, for example, Xu and Taylor (1995)], and hence we adopt it here. We estimate the GARCH parameters under the assumption that the errors from the mean equations are conditional normal.[29] The starting values for the estimations are obtained by OLS, and the quasi-

over the same period.

[29] We tried a number of explanatory variables in the mean equation to identify significant ones for inclusion in the estimations. The candidate variables that we experimented with are lagged returns (to remove any autoregressive process), day of the week, month, domestic (U.S.), and foreign interest rates. Tests of these variables are insignificant and hence they are not included in the final estimates. This is consistent with many exchange rate forecasting models which do not include conditional mean dynamics [see, for example, Diebold and Nason (1990)].

Newton nonlinear optimization method (the BFGS) is employed to estimate the parameters and obtain reliable estimates of the standard errors. The GARCH results are reported in Table 3.4. The low estimated standard errors of the variance equations (significant variance equation parameters) support the GARCH methodology. In the variance equations, the value of $\alpha_1 + \beta_1$ is close to unity, which indicates possible existence of an integrated GARCH (IGARCH) process [see Engle and Bollerslev (1986)].[30] However, consistent with many other studies [see, for example, Baillie and Bollerslev (1989)], we did not employ this technique here and we allow free estimates of the coefficients.

The GARCH equations for all the currencies are strikingly similar. The estimated constant in the mean equation is insignificant for all the currencies at 5%. This is consistent with most of the literature on exchange rate predictions, which suggest that there is no evidence of significant exchange rate appreciations or depreciations of major currencies over long periods. It should be pointed out that although the BP shows significant first order autocorrelation (see Table 3.3), this relationship disappears in the GARCH setup and, hence, we dropped the lagged return from the mean equation. The parameters of the variance equations are all significant which suggest a non-trivial dynamic structure in the second order moments. The persistence parameter estimates dominate the variance equations across the currencies. They are high and positive for all the currencies which suggest volatility clustering of the exchange rate returns. A notable surprise is that the Box Pierce statistic is insignificant for the JY (see Table 3.3), but its variance equation display strong persistence.

[30] JP Morgan RiskMetrics corresponds to IGARCH model without a drift.

We test the standardized residuals from the mean equation to assess the validity of the conditional normal assumption used in estimating the parameters. The summary statistics of the standardized residuals are reported in Table 3.5. The table shows that the effect of the GARCH filtering is not uniform on the kurtosis across the currencies. The kurtosis of the BP and the DM decreased while it increased for the CD, the FF and the JY. Decreases in kurtosis of the BP and the DM after the GARCH filtering reflect the GARCH hypothesis that part of the stochastic component of the exchange rate returns is due to different conditional normal distributions (removing the heteroskedasticity associated with mixtures of normals will result in a reduction in the sample kurtosis). On the other hand, the increases in kurtosis of the CD, the FF, and the JY suggest that the GARCH filtering was not effective in removing some of the kurtosis as expected. This would occur if there are extreme shocks to the exchange rates. The GARCH filter is unable to reduce the effect of these extreme shocks because volatility adjustments lagged the shocks by one period. Thus, in relatively tranquil periods, sudden shocks to the exchange rate would result in large standardized residuals. This emphasizes the point that VaR models are often ineffective when the market behaves abnormally because they have no power to predict sudden shocks. The coefficients of skewness decreased for the BP and the DM but increased in absolute terms for the remaining currencies.[31] Table 3.5 also shows the standardized residual values corresponding to the 1^{st}, 5^{th}, 95^{th} and 99^{th} percentiles. These values suggest that the distributions of the standardized residuals are not, in general, fat-tailed at the 95% relative to the Guassian benchmark, but they are at

[31] If the standardized residuals are iid, then the conditional kurtosis should not be greater than the unconditional kurtosis. Likewise, the conditional skewness should not be greater than the unconditional skewness [see, for example, Engle and González (1991)]. The fact that some of the exchange rates violate

99% [only one value, 2.21 for JY, is lower than the 1% normal value of 2.33]. It is also worth noting that the extreme values of the standardized residuals increased relative to the raw data, which explains the increase in kurtosis of some of the exchange rates after running them through the GARCH filter.

Both the chi-square and the Jack-Bera tests reject the null hypothesis that the exchange rate returns are conditionally normal (although the test statistics are lower than their unconditional counterparts).[32] Since the standardized residuals are not normally distributed, the distributional assumption underlying the parameter estimates is inappropriate and hence the empirical approach and the conventional normal VaR calculations will, in general, be different.

3.5.1.2. Forecasting

The real test of a risk management methodology is out-of-sample performance. The risk manager, by definition, obtains VaR estimates in real time and hence must use parameters obtained from an already observed sample in order to evaluate the risks associated with current and future random movements in the risk factors. Hence, a true test for the different VaR methods is their performances outside the sample used to estimate the underlying parameters. This study, therefore, focuses solely on assessing the VaR estimates using out-of-sample forecasts.

To measure the out-of-sample performance, we focus on a one day ahead forecasts since VaRs are usually calculated for this time horizon for internal

---

these conditions is an indication that the standardized residuals are not iid. This could be due to misspecified model or distribution.

[32] This is to be expected as the GARCH parameters are chosen such that the innovations fit the normal likelihood as much as possible given the sample.

management. For these one-day ahead forecasts, we use the estimated variance equation to forecast the variance for the next day. Conditional on this variance, the next day's risk structure is normal under the GARCH hypothesis and, hence, the conditional normal distribution is used to calculate the VaRs at 99% and 95% confidence levels for both long and short positions in the U.S. dollar. Under the Barone et al. approach, we draw 5000 times (with replacements) from the standardized residuals and multiply these draws by the estimated standard deviation for the next day. These scaled standardized residuals are used to forecast returns to construct the empirical distribution of the exchange rate risk for the next day [see Barone et al. (1999) for detail description of this procedure]. After calculating the VaR, the sample is rolled forward one day and the forecasting procedure is repeated to predict the next day's VaR for the different approaches.[33]

The evaluation criteria, discuss earlier (section 3.3), are used to evaluate the performance of the different modeling techniques. The quantitative precisions of the VaRs obtained using the normal and the empirical distributions to forecast the future distributions of the exchange rate returns are shown in Table 3.6. The performance criteria shown in the table include the number of exceedences in the last 250 days at 99% confidence, the exceedence ratio, the degree of exceedence, mean absolute deviation (MAD) and the root mean square error (RMSE). MAD measures the mean absolute deviation of the VaR estimates from the actual exchange rates. RMSE is calculated as the square root of the squared deviations of the VaR from the actual exchange rate. These measures would be used as supplementary yardsticks in assessing the performance of the various VaR models since they focus on performance at the center of the distributions as

[33] This is different from using a static sample where the parameter estimates from the sample period are used throughout the forecasting period as in Christoffersen (1998). The methodology adopted here is more

oppose to the tails, which is relevant in predicting VaR. The smaller the MAD and the RMSE, the better the model predicts the exchange rates.

The Bank for International Settlement (BIS) recommends that a model is in the green zone if in the last 250 days the number of losses in excess of the model's values at risk at 99% confidence is *not* greater than four. The losses in excess of the VaRs at the 99% confidence for the different models in the last 250 days are reported in brackets in Table 3.6. From the results, the standard conditional normal distribution model fails the green zone requirement on long positions in the U.S. dollar for the GM and FF. The model, however, passes the green zone test for short positions in the U.S. dollar across all the currencies. The simulated empirical distribution approach, on the other hand, meets the BIS green zone stipulation for both short and long position on the U.S. dollar across all currencies. In addition, the number of exceedences obtained using the empirical distribution forecasts in the last 250 days is not greater than those under the normal distribution forecasts across all currencies for both shorts and longs. Thus, based on the BIS criterion, the empirical approach performs better than the standard GARCH method.

The BIS recommendation, however, does not penalize models that overestimate risk (low exceedence counts) and, hence, cannot appropriately evaluate VaR models. Exceedence ratios overcome this problem, and punish both overestimation and underestimation of risk. An exceedence ratio greater than one suggests that the model underestimates risk and a ratio less than one indicates that the model overestimates risk. The exceedence ratios (see Table 3.6) suggest that the forecasts of the standard GARCH/normal procedure overestimate the risk at 95% level of confidence by 11% on short positions in the US dollar and by 12% on long positions for the JY. The model

---

realistic as institutions would continuously update their database to calculate their VaRs.

under-predicts the risk of the GM and the CD by 10% and 12%, respectively, on short positions in the US dollar. With the exception of the above, the model did not over- or under-estimate any position in the US dollar by more than 7%. Thus, in general, the model performs well at the 95% confidence level on both short and long positions in the US dollar across the currencies.

The empirical forecasting method overestimates the JY exchange rate risk by 11% on long positions, but underestimates the risk on short positions by 24%. It also under estimates the risk of the CD by 14% on short positions. For the remaining positions, the empirical forecasts do not over- or under-estimate risk by more than 8%. In general, the violations of VaR under the two approaches (normal and empirical distributions' forecasts) are comparable at the 95% confidence on both short and long U.S. dollars across the exchange rates, with the exception of the JY. The JY results are interesting. While the normal method over predicts risk by 11% on shorts, the empirical approach underestimates risk by 24% on the same position.

The observed results can be explained, in part, by reference to some of the empirical characteristics of the exchange rate returns. The average critical values of the GARCH filtered standardized residuals over the forecast period support the sample period evidence that the distributions of the standardized residuals are not fat-tailed at 95% for both long and short positions in the US dollar. The 90% confidence bands are [-1.63, 1.65], [-1.70, 1.65], [-1.50, 1.66], [-1.65, 1.63], and [-1.62, 1.46] for the BP, GM, CD, FF and the JY, respectively (see Table 3.5). The average empirical interval forecasts of the standardized residuals are [-1.60, 1.63], [-1.66, 1.64], [-1.58, 1.62], [-1.61, 1.64], and [-1.60, 1.44] for the BP, GM, CD, FF and JY, respectively. These empirical bands (both from the sample and the forecast periods) suggest that the normal distribution forecast is sufficient to account for tail-fatness at the 95% confidence for both short and long positions across most of the currencies. Thus, it is not surprising that the normal and the empirical forecasts do not produce significantly different results at the 95% confidence across the exchange rates, with the exception of the JY. The empirical bands suggest that the normal distribution is 'too' fat tailed at 95% confidence for short positions in the U.S. dollar for the JY and, therefore, explains why the normal approach over predicts the risk for this position. The information from the 90% confidence intervals of the standardized residuals do not provide a direct explanation of why the empirical approach under forecasts the JY exchange rate risk.[34]

The relationship between the empirical sample interval and the average forecast interval can be used to shed some light on the observed performance of the VaR estimates under the empirical approach. Compared to the sample band, the forecast bands at 90%, in general, contract slightly. In the case of the JY, a shrinking band (lower VaR forecasts) would result if the forecast distributions of the standardized residuals, on average, gain density mass in the interval [-1.62, 1.46] which is the 90% sample confidence interval. This would occur if the forecast variances are, on average, less than the mean sample period variance.[35] The forecast variances would be less than the average sample variance if new information arriving suggests that the future is less volatile (less exchange rate shocks) and/or the variance equation is not responsive enough to the exchange rate shocks. While the former would not result in under forecasting, the latter

[34] If the 95th percentile value of the standardized residuals were increasing through the forecast period, then using the lower sample values will, in general, result in under forecasting the risk. The 95th percentile sample value of 1.46 is higher than the average value over the forecast period (1.44), and, hence, does not provide an explanation for the under forecasting.

would. An inspection of the JY variance equation suggests that the forecast variances are not very sensitive to the exchange rate shocks (its shock parameter estimate is the least among the currencies) which, possibly, explains why the empirical approach under forecasts the JY exchange rate risk. This captures a conflict between the distributional assumption used to estimate the parameters and the one used to forecast the future risk structure. Since the empirical distribution of the JY exchange rate is not fat tailed relative to the Gaussian distribution, the parameters are estimated such that the normal distribution would not over predict the JY exchange rate risk as much as possible. In other words, the normal distribution has enough density at the tails to forecast the exchange rate risk and, hence, high time variation in the conditional variances is not required. This would result in a low estimate of the exchange rate shock parameter in the variance equation.[36] While this is consistent with the normal distributional hypothesis it poses problems for the empirical distribution.

At the 99% confidence, the normal distribution forecasts under predict the risks of all the exchange rates for both longs and short positions in the U.S. dollar. The underestimation is especially severe on long U.S. dollar positions for all the currencies, ranging from 63% for the CD to 103% for the JY. The empirical approach, on the other hand, did quite well in forecasting the VaR with the exceptions of the shorts in U.S. dollars for the CD and on both positions for the JY. The empirical forecasts overestimate the CD dollar exchange rate risk by 22% on short U.S. dollar positions, and they underestimate the JY exchange rate risk by 20% and 32% on long and short positions in

[35] Note that the sampled standardized errors are multiplied by the forecast standard deviations to adjust them to the current volatility condition.

the U.S. dollar, respectively. The results clearly show that the empirical approach dominates the normal GARCH technique at 99% confidence. The reason for this can be unraveled by reference to the empirical characteristics of the exchange rate returns. The 99% confidence bands for the sample period are [-2.69, 2.45], [-2.72, 2.54], [-2.56, 2.65], [-2.64, 2.52], and [-2.97, 2.21] for the BP, GM, CD, FF and the JY, respectively (see Table 3.5). The average empirical interval forecasts of the standardized residuals are [-2.64, 2.49], [-2.75, 2.57], [-2.64, 2.65], [-2.62, 2.55], and [-3.09, 2.27] for the BP, GM, CD, FF and JY, respectively. With the exception of short positions in the U.S. dollars for the JY, the empirical percentiles from the sample and the forecast periods suggest that the normal distribution does not have enough mass at the tails to properly forecast VaR at 99% confidence for both short and long positions across the currencies. In particular, the empirical percentiles suggest that the normal distribution will underestimate the exchange rate risks at 99% confidence on both positions, and the empirical results confirm this across the currencies. Both the sample empirical standardized residuals and the forecast standardized residuals suggest that the normal distribution would overestimate risk on short position in the U.S. dollar for the JY. The result, however, indicates that the normal forecast underestimates risk by 14%.

An important lesson from these results is that the empirical approach dominates the standard normal procedure when the percentile of the GARCH filtered innovations is fat-tailed at the relevant confidence level relative to the Gaussian benchmark. Indeed, Duffie and Pan (1997) point out that if one is concerned with measuring the VaR of direct

---

[36] The likelihood function is maximized with respect to the entire sample. Thus, this type of result (low shock parameter estimate) may not be obtained even when the normal is too fat at the 95% confidence since the parameter estimates depends on the empirical distribution of the entire sample.

exposures then the pertinent measure of tail fatness is the number of standard deviations represented by the associated critical value. The results, therefore, suggest that benefits can be derive from using the empirical distribution methodology if the distribution underlying the parameter estimates does not have enough density at the tails to account for tail fatness at the relevant confidence level. Thus, the empirical approach is superior to the standard procedure when the GARCH distributional assumption is inappropriate at the relevant percentile. Inappropriate likelihood functions would produce unsuitable estimates of the GARCH parameters. One way to address this theoretical concern, at least partially, is to invoke the Quasi Maximum Likelihood Estimation (QMLE) argument. The QMLE procedure does not change the normal distribution likelihood but calculates more robust estimates of the covariance parameter estimates. Under fairly weak conditions [see Bollerslev and Wooldridge (1992)] the resulting estimates are consistent even when the conditional distribution of the GARCH residuals is non-normal. Note that the QMLE technique produces the same parameter estimates as the standard MLE procedure, and, therefore, the two procedures will generate the same VaR estimates.[37]

Apart from the above quantitative discussion, we use the unconditional, the independence, and the conditional likelihood ratio tests [see, Christoffersen (1998)] to formally test the accuracies of the models. These tests are reported in Table 3.7. The unconditional test corresponds to the standard test employed by many researchers [for example, Baillie and Bollerslev (1992) and McNees (1995)], which tests whether the observed exceedence is significantly different from the expected exceedence. This test essentially captures the appropriateness of the distribution used to forecast the

[37] Different results will be obtained only when the higher standard errors obtained under the QMLE causes the researcher to drop some of the parameters from the model.

innovations of the exchange rate returns. The test, however, is unable to determine whether the observed exceedences occur randomly over time or cluster in some fashion. The independence test solves this problem by examining the randomness of the observed violations over time. This test, therefore, captures how effective the higher order dynamics of the exchange rate return series is modeled. The conditional test combines both the coverage and the independence tests and, hence, examines the appropriateness of distributional assumption as well as the variance dynamic structure.

The number of observations in the forecast period is not large enough to calculate the independence and conditional test statistics at 99% confidence for some of the exchange rates and we report these results as N/A. This is because observed consecutive exceedences of VaRs are zero for some of the exchange rates.[38] This shows the difficulty of empirically validating models at very high confidence levels such as the 99% level commonly used by various institutions and regulatory bodies. The results indicate that the tests do not reject the VaR models at 95% level of confidence on either the short or the long position in the US dollars for both methods (normal and empirical distributions) at 10% significance level of test. Thus, at the 95% confidence, the tests do not reject either model across the currencies. This is consistent with the earlier discussion that both models perform reasonably well across the currencies at the 95% confidence level. A case of interest is that the JY exchange rate returns pass the test at 95% level of confidence on short U.S. dollar positions although it under forecasts the risk by 24% (it has the highest likelihood ratio statistic at the 95% confidence level for the unconditional tests).

[38] The expected number of consecutive exceedences at 99% confidence is .166.

At 99% level of confidence, however, the empirical approach is superior. The results reject the normal distribution forecasts for the GM, the FF, and JY on long positions in the U.S. dollar. The results show that the rejection is due to improper distributional assumption rather than inaccurate modeling of the variance dynamics. This is consistent with the empirical evidence that the normal distribution is too thin at 99% confidence level to model the exchange rate returns. Christoffersen also finds that the empirical approach may be superior at higher levels of confidence (in his study, the empirical approach is superior at 95% confidence interval, that is 97.5% on short and long positions).[39]

3.5.2. The t-distribution Likelihood

The GARCH estimates obtained using the t-distribution likelihoods are reported in Table 3.8. We use two t-distribution likelihoods to estimate the GARCH parameters. In the first one, we choose the degree of freedom (selected degree of freedom) such that the chosen t-distribution approximates the empirical percentile at 99% for both short and long positions as much as possible and, in the second, we allow RATS to estimate the degree of freedom (estimated degree of freedom) along with the other parameters.[40] The estimated degrees of freedom are lower than the selected degrees of freedom across the currencies. This allows us to accomplish two objectives. First, we are able to evaluate the performance of VaR estimates when the degree of freedom is selected so that the t-distribution matches a particular empirical percentile. Second, the lower estimated

[39] Christoffersen (1998) use a static sample period percentile as the empirical forecast.
[40] In choosing the selected degree of freedom, we first find the mean of the absolute values of the 1st and 99th percentile of the standardized historical exchange rates. We then select a degree of freedom parameter

degrees of freedom allow us to obtain richer insights on the performance of the GARCH models as the innovations are progressively assumed to come from distributions with fat tails.

3.5.2.1. Estimation

From Table 3.8, the GARCH parameter estimates are, in general, similar to the ones obtained using the normal distributional assumption. Like their normal counterparts, the constant terms in the mean equations are all insignificant at 5% level of test, with the exception of the BP when the degree of freedom is estimated. The variance equations are also dominated by the persistence parameters, and they suggest that the volatilities of the exchange rates cluster like their normal likelihood analogs. The estimated coefficients of the shocks are also significant across the exchange rates for the two t-distributions.

The effects of changing the distributional assumption on the parameter estimates of the variance equations are not uniform across the currencies. However, for each currency, we can broadly observe some systematic differences between the t-distributions and the normal likelihoods' estimates. The absolute values of the estimated constant terms in the variance equations increased for the BP and the FF, but decreased in the remaining exchange rates compared to their normal likelihood counterparts. This suggests that, relative to the normal estimates, the estimated constant terms would increase the time variation in the forecast variances of the BP and the FF, but decrease those of the remaining currencies. The estimated coefficients of the shock and the lag variance increased for the CD under the t-distributional assumptions. This suggests that,

---

such that the 99% percentile theoretical value is as close as possible to the empirical average. The selected degree of freedom is then held constant throughout the forecasting period.

for the CD, exchange rate shocks are more influential on the variance of the exchange rate returns and their effects on the forecast variances persist for longer periods than the normal distribution estimates suggest. The estimates of the shock coefficients of the FF and the BP increased, but their persistence parameter estimates decreased relative to their counterparts under the normal hypothesis. These suggest that, although, the exchange rate shocks increase the forecast variances more than the normal estimate indicates, the increases in the variances are less persistent. The shock coefficient estimate of the forecast variance decreased for the DM, but the persistence parameter estimate increased. This suggests that shocks have less influence on the predicted variances under the t-distribution hypothesis, but the effects are more persistent than the normal hypothesis indicates. The JY is the only currency for which the t-distribution likelihoods did not induce systematic effects on the parameter estimates. While the two t-distribution likelihoods produced lower persistence parameter estimates, the shock coefficient estimate decreases when the degree of freedom is selected but increases when RATS estimates the degree of freedom. Thus, compared to the normal hypothesis, shocks have less influence on the predicted variances when the selected degree of freedom t-distribution is used, but have more influence when the degree of freedom is estimated to maximize the likelihood function. The shock effects are, however, less persistent in both cases relative to what the normal distribution estimate suggests.

There are also differences in the estimates of the two t-distribution likelihoods. In general, the fatter tailed t-distributions tend to emphasize the deviation from the normal estimates. The exceptions are the estimate of the constant term in the variance equation of the DM, the constant and the persistent parameter estimates of the FF and all the estimates of the variance parameters of the JY. The quantitative differences in the estimates of the variance parameters will generate different time variations in the variances and, hence, kurtosis associated with stochastic volatility would be different for the different distributional assumptions. The estimated variances will exhibit more time variation the higher the constant term and/or the higher the shock coefficient estimates. The effects of changes in the estimated persistence parameters on the conditional variances are not monotonic and, consequently, are not predictable *a priori*.

We report the various statistics of the standardized residuals in Table 3.9. These statistics allow us to compare the effects of the GARCH/t-distribution filters on the exchange rates and also examine the validity of the t-distributions. The kurtosis of the residuals of the BP and the DM decreased relative to their raw values (see Table 3.1.) but those of the CD, the FF and the JY increased for both t-distributions (the GARCH/normal filters had similar effects on the kurtosis of the exchange rates).[41] Thus, time varying variances decreased the kurtosis associated with the BP and the DM, but increased the kurtosis associated with the remaining currencies. As discussed in section 3.5.1.1, an increase in kurtosis would occur if there are sudden extreme shocks to the exchange rates. This is because the conditional variances respond to exchange rate shocks with a lag. Thus, in tranquil period when the variances are low, a sudden shock would result in a high standardized residual.

The kurtosis of the errors can be compared to their theoretical counterparts to determine whether the excess kurtosis in the exchange rate returns are sufficiently represented by the t-distributions. Given the estimated degree of freedom ($\hat{f}$), the

[41] Engle and González-Rivera (1991) show that if the errors are iid, then the conditional kurtosis will not be greater than the unconditional kurtosis.

theoretical conditional kurtosis is equal to $3(\hat{f}-2)/(\hat{f}-4)$, where $\hat{f} > 4$ [see Bollerslev (1987) and Kendell and Stuart (1969)]. This can be compared to the sample analog of the standardized residuals, calculated as the mean of $\hat{e}_t^4 / \hat{h}_t^2$, for t=1,...,s. The first order approximation of the asymptotic variance of the estimated conditional kurtosis is given by $36\text{var}(\hat{f})/[1-4/\hat{f}]^4$. However, the behaviour of the likelihood function for small estimated degrees of freedom is unclear [Baillie and Bollerslev (1989)] and hence the standard errors of $3(\hat{f}-2)/(\hat{f}-4)$ are not reported. Informal examination of the standardized residual, however, indicates that the CD, the FF, and the JY exhibit unaccounted for kurtosis when the degree of freedom is selected. The residuals from the estimated t-distributions yield sample kurtosis that are all less than their theoretical values and, therefore, suggest that the estimated t-distributions over-account for the empirical kurtosis. Thus, whereas the selected t-distributions did not sufficiently account for the empirical kurtosis for some of the exchange rates and, hence, would under predict the risks of these currencies, the estimated t-distributions over account for the empirical kurtosis of all the currencies and, hence, would over predict the risks associated with the exchange rates.

A striking observation is that the t-distributions that maximize the likelihood functions are fat tailed relative to the distribution of the standardized residuals and, thus, would potentially yield poor VaR forecasts. This highlights the conflict between maximizing the likelihood function, which is concerned with the entire distribution, and VaR forecasts which depends on modeling the tails accurately. Since the likelihood estimations depend on the whole distribution, fitting the shoulders and the centers of the distributions are also important.[42] This suggests that the relatively fat tailed t-distributions increase the likelihood functions because the deteriorations in the tails are overcompensated for by the improvements in the centers and the shoulders.[43] This assertion can view graphically from figure 3.1. The figure shows the distributions of the residuals from the selected and the estimated t-distributions. It can be seen that there is not much difference between the distributions of the GARCH/selected-t and the GARCH/estimated-t residuals at the tails, and the differences are primarily at the shoulders and the center. This supports the argument that, given the class of t-distributions, the fat tailed t-distributions fit the data better mainly because they capture shoulders and the centers better.

The 95% percentile values of the standardized residuals of the t-distributions suggest that the residuals are not fat-tailed relative to their theoretical benchmarks for both t-distributions across the currencies. The 95th percentile theoretical values of the estimated t-distributions are, particularly, higher than their sample counterparts and, hence, suggest that these t-distributions over account for the empirical densities at the 95% confidence level. At 99% level of confidence, the selected t-distribution values lie between the absolute values of their sample counterparts. To a large extent, this is because the degrees of freedoms were selected so that the 98% central interval values are as close as possible to their empirical analogs. The estimated t-distributions have theoretical values that are higher than their sample counterparts, which indicate that these

[42] Although, the data points in the centers and the shoulders are less weighted in the likelihood function, they can exert strong influence when they are many.
[43] The likelihood functions of the estimated t-distributions are higher than those of the selected t-distributions across the currencies.

distributions are also "too" fat at the 99% level of confidence to properly predict the values at risk of the exchange rates.

Relative to the standardized residuals from the normal hypothesis, the extreme values from the selected t-distributions increased in absolute terms across the currencies, but the values from the estimated t-distributions do not portray any particular systematic pattern. For the estimated t-distributions, the extreme values (in absolute terms) increase for the BP and the CD, but decrease for the FF. In case of the DM, the maximum value decreases while the minimum value increases. It is also worth noting that the extreme values of the standardized residuals for both t-distributions increase relative to their raw values, which explains the increase in kurtosis of some of the exchange rates after running them through the GARCH-t filters.

3.5.2.2. Forecasting

The forecasting results from using the selected t-distributions and the estimated t-distributions are reported in Tables 3.10 and 3.11, respectively. From the tables, the exceedences in the last 250 days at 99% confidence level (reported in brackets) are not greater than four for both t-distribution likelihoods across the currencies. This suggests that both the empirical and the t-distributions' forecasts of the values at risk of the exchange rates pass the BIS green zone test. For the selected degrees of freedoms, the exceedences in the last 250 days of the theoretical forecasts are lower on short positions for the BP and the JY, but are higher on long positions for the GM and the FF relative to the empirical forecasts. Thus, in terms of lower violation counts, the empirical forecasts do not dominate the selected t-distribution forecasts. The estimated t-distributions, on the other hand, produce forecast violations that are less than the empirical VaR violations across all currencies in the last 250 days. This suggests that the estimated t-distribution forecasts are superior since the BIS criterion does not penalize overestimations. Thus, based on the BIS recommendation, the results here suggest that the forecasts of the two types of t-distributions, unlike the normal distribution forecasts, are not outperformed by their empirical counterparts.

The exceedence ratios suggest that the t-distributions tend to over predict the risk associated with the exchange rates for both short and long positions in the U.S. dollar at 95% level of confidence. As is expected, the overestimation of risk is more pronounced when the degree of freedom is estimated. The maximum exceedence ratio of the estimated t-distribution forecasts is only 0.57. The overestimation of risk by the two t-distributions at 95% confidence level is consistent with the empirical evidence that the distributions of the exchange rate returns are not fat-tailed at this level relative to the theoretical benchmarks on both short and long positions for all the currencies. The selected t-distribution likelihoods combined with the empirical distribution forecasts yield VaR violations that are within 10% of their expected values for all the currencies, except short positions in the U.S. dollar for the JY and both positions for the CD. Specifically, the procedure under predicts the CD exchange rate risk by 12% on long positions in the U.S. dollar and by 20% on shorts. For the JY, the methodology under forecasts long positions in the U.S. dollar by 29%. Using the estimated t-distribution likelihoods, the empirical forecasts produced violations that are in excess of 10% of the expected on both positions for the CD and the JY. In particular, this procedure under predicts the risks of long positions in the U.S. dollar by 13% and shorts by 17% for the

CD. In case of the JY, the methodology over forecasts long positions in the U.S. dollars by 12% and under predicts short positions by 27%.

At 99% level of confidence, the selected t-distributions over predict the risks of some of the exchange rates, while they under estimate the risks of others. In particular, the exchange rate risks associated with long positions in the U.S. dollar are over predicted for the GM and the JY and, also, on both positions for the FF (exceedence ratios are above 1.1). On the other hand, the selected t-distributions under predict the risks of short positions in the U.S. dollar for the BP and the JY (exceedence ratios are below 0.9). At the same confidence level, the forecasts of the estimated t-distributions are all lower than their expected values. Thus, as expected, the estimated t-distributions are overcautious in predicting the exchange rate risks across all currencies and the maximum exceedence ratio is only 0.54.[44] Again, the over estimation is consistent with the empirical evidence that the estimated t-distributions are too fat at the tails to properly model VaR at the 99% confidence level across the currencies.[45] It is worth noting that the exceedence ratios obtained using the t-distributions are all less than their normal counterparts, which reflect the fact that the t-distributions are fatter at the tails than the normal distribution. At the 99% level of confidence, the empirical forecasts generated with the selected t-distribution likelihoods achieve exceedences within 10% of their expected values with the exceptions of long positions in the U.S. dollar for the BP and of both positions for the JY. The model under projects the risks of the BP and the JY exchange rates by 14% for long positions in the U.S. dollar, and under predicts the risk

[44] Christoffersen (1998) also finds that the GARCH/t-distribution over predicts risk of the BP, GM, JY and the Swiss Franc at 97.5% confidence on both longs and shorts.

[45] The results are also consistent with the empirical evidence that the kurtoses in the errors are over accounted for by the estimated t-distributions across the currencies.

of short positions by 27% for the JY. The empirical forecasts generated with the estimated t-distribution likelihoods also achieve exceedences within 10% of their expected values for all the currencies, with the exceptions of long positions in the U.S. dollar for the CD and shorts for the JY. The VaR forecasts under predict the risk of long positions in the U.S. dollar by 14% for the CD and by 33% for short positions in the U.S. dollar for the JY.

At this point, it is worthwhile to note a few comparative observations from the above discussions. The selected t-distribution VaR forecasts outperform their estimated t-distribution counterparts at both 95% and the 99% confidence levels. This is because the estimated t-distributions are too cautious in predicting the exchange rate risks at both confidence levels. It over estimates the risks of all the currencies across both positions at the 95% and 99% levels of confidence. The exceedence ratios from the estimated t-distribution forecasts range from a low of zero to a high of only 0.57. The reason is that maximization of the likelihood function deals with the entire distribution while VaR forecasts depend on the tails of the distribution. This also explains why the empirical distribution VaR estimates dominate the estimated t-distribution forecasts. Thus, unlike the normal distribution, the empirical distribution forecasts dominate the estimated t-distribution VaR forecasts because of over prediction of risk and not under prediction at the 99% level of confidence.

The empirical distribution forecasts do not dominate the selected t-distribution forecasts at both levels of confidence. At the 95% confidence level, the empirical approach outperformed the selected t-distribution forecasts on long positions in the U.S. dollar for the BP exchange rate, but it did not dominate the t-distributions VaR forecasts for the remaining currencies. For these remaining currencies, the empirical distribution forecasts generally outperformed the selected t-distribution forecasts on the long positions in the U.S. dollar while the selected t-distribution forecasts are superior on the short positions. Thus, unlike the estimated t-distribution forecasts, the empirical forecasts are not clearly superior to the selected t-distribution forecasts at 95%. At the 99% level of confidence, the empirical forecasts dominated the selected t-distribution forecasts on long positions in the U.S. dollar for the GM and the shorts for FF, but its forecasts are not dominant in predicting the remaining exchange rates. For the remaining currencies, the empirical distribution VaR forecasts of the BP, CD, and the JY are superior on shorts positions in the U.S. dollar, while the selected t-distribution forecasts are better on the longs. Thus, unlike the normal distribution forecasts, the empirical distribution forecasts do not dominate the theoretical forecasts of the selected t-distribution across the currencies at 99% level of confidence.

A more formal test based on Christoffersen (1998) is used to examine the performance of the models and shed light on violations that are due to wrong distributional assumption and those that are due to improper modeling of the variance dynamics. The results for the selected t-distributions are reported in Table 3.12. At 95% confidence, only the independence test of the BP exchange rate risk of long positions in the U.S. dollar fails the likelihood ratio test at 5% level of significance. This suggests that, although the t-distributional assumption is appropriate, the variance dynamics is inaccurate, which results in the failure of the entire model at 10% level of significance. Also, the empirical forecasts generated by the selected t-distribution forecasts fail the distributional test at 10% level of significance on short positions in U.S. dollars for the

JY. At 99% level of confidence, all the calculable likelihood ratio statistics are lower than the critical value at 5% and, hence, the tests do not reject either the selected t-distribution or the empirical distribution forecasts. However, for the JY exchange rate risk, the selected t-distribution forecasts of long positions in the U.S. dollar just marginally passed the distributional assumption test at 5% level of significance.

The unconditional tests do not indicate that the forecasts of the empirical distribution are superior to the selected t-distribution forecasts. At 95% confidence, the selected t-distribution forecasts dominated the empirical distribution ones (has lower likelihood ratio statistics) on short positions in the U.S. dollar for the GM, CD, JY exchange rate risks, and on long positions for the CD and the FF. At 99% level, the selected degrees of freedom t-distribution estimates are superior in forecasting the VaRs of long U.S. dollar for the BP and the CD, and of shorts for the GM and the JY. These unconditional likelihood ratio tests support the earlier informal discussions that the empirical distributions' VaRs did not dominate the selected degrees of freedom t-distribution VaRs across the currencies. Thus, unlike the normal distribution, the t-distribution is not overly confident and hence does not underestimate the risk of the exchange rates across the currencies at 99%. Indeed, one can argue that the selected t-distribution would achieve better results than it did in this study if they are selected to match the empirical values of specific positions rather than the average of both short and long position as we did here.

The likelihood ratio statistics of the estimated t-distributions are reported in Table 3.13. The results show that the estimated t-distributional assumption is inappropriate for forecasting all the exchange rate returns at both 95% and 99% confidence levels. This is confirmed by the likelihood ratio tests, which show that variance dynamics are properly modeled and, hence, the failure of the model is due to the poor distributional assumption. These results are consistent with the observation that the estimated t-distributions are too fat at the tails to properly model VaR at both confidence levels. The empirical forecasts, on the other hand, passed all the distributional tests. This approach, however, fails the independence test in the forecasting the BP exchange rate risk on long positions in U.S. dollars at 95% confidence level. The implication of this is that while the empirical forecast distribution is appropriate, the variance dynamic structure is incorrect. Since the same estimated variance dynamics model passed the independence test when the theoretical distribution is used to forecast the VaR, the failure does point to a possible inappropriateness of combining the estimated variance equations with the empirical forecast distributions. The variance parameters are estimated so that the distribution of the standardized residuals mimics the estimated t-distribution as much as possible. As the unconditional likelihood ratio test statistics show, the empirical distributions are different from the t-distributions (the theoretical distributions are rejected while the empirical distributions are not). From the VaR perspective, this distributional difference arises because the estimated t-distributions are fatter at the tails than the distribution of the standardized residuals. Thus, given the fat tailed theoretical distribution, the estimated variance dynamics is sufficient to ensure that the VaR violations do not exhibit excessive dynamic clusterings. This, however, is not sufficient for the thinner tailed empirical distributions. Thus, increasing time variation in the conditional variances would produce better results given the distribution of the standardized errors. However, increasing time variation may lead to a failure of the independence test of the theoretical distribution forecasts and, hence, this demonstrates a possible conflict between modeling time variation to suit the theoretical distribution and modeling it to suit the empirical distribution of the standardized residuals.

The explanation of the cause of the failure of the dynamic forecasts of the BP exchange rate warrants more empirical information to fully understand the source of the problem. In particular, the 5th and the 95th percentile values of the adjusted standardized errors over the forecast period and the amount of time variation in the conditional variances will help shed some light on the sources of the failure. In general, dynamic predictive failure can occur either because there is not enough time variation in the conditional variances or there is too much time variation, given the assumed forecast distribution.[46] In case of the BP, the failure of the independence test is due to many consecutive violations of the VaR forecasts (12 consecutive violations compare with an expected value of 4.15). This suggests that the estimated variance equations are inappropriate because the time variation in the conditional variances is not sufficient, given the forecast empirical distributions. Although, the estimates of the variance equations can give us some idea about the amount of time variation in the forecast variances, we did not depend on these estimates to infer the volatility in the conditional variances for two reasons. First, the change in the parameter estimates caused by the different distributional assumptions did not change the time variation in the conditional variances in a particular direction, and the effect of changes in the persistent parameter estimates is ambiguous (see section 3.5.2.1). Second, we use a rolling sample and hence the parameter estimates are not static (although they do not display wide variations). In view of these, we calculate the time variation in the projected variances directly by calculating their standard deviations. The standard deviations of the projected variances of the different exchange rates over the forecasting period are reported in Table 3.14.[47] These standard deviations estimate the volatility associated with the forecast variances and, hence, they capture, to some degree, the volatility in the exchange rates associated with time-varying variances.

The results from the table show that, for the BP exchange rate, the estimated t-distribution likelihood produces the highest time variation in the conditional variances compared to the normal and the selected t-distribution likelihoods. This is counterintuitive as, *ceteris paribus*, one will expect the estimated t-distribution to generate the least time variation.[48] An examination of the average empirical forecast intervals, which we report in Table 3.15, sheds some light on this anomaly. From the table, the empirical forecasts generated by the estimated t-distribution likelihoods surprisingly produce the smallest width at both 90% and 98% central confidence relative to the normal and the selected t-distribution for the BP. This emphasizes the point that, in terms of maximizing the likelihood function, fitting the distribution at the center and the shoulders are also every crucial in determining the estimated variance equation. Thus, the relatively high time variation resulting from the estimated t-distribution likelihood is mainly to model the center and the shoulders of the distribution, and it is not sufficient to

---

[46] Inadequate time variation in the conditional variances will lead to more than expected number of clustering of VaR exceedences while too much variation will lead to less than expected clustering of violations.

[47] Note that the time variation in the forecast variances does not depend on the methodology used to forecast the future risk structure and, therefore, is the same for both the empirical and the theoretical forecasts with the same underlying likelihood function.

[48] This is because we expect the estimated t-distribution to produce the fattest densities at the tails of its standardized residuals and, hence, should require less time variation in the conditional variances to account for the empirical kurtosis.

compensate for the lower empirical densities it causes at the tails to properly forecast VaR. This suggests that, whereas time variation in the forecast variances is sufficient for the fat-tailed theoretical distribution, it is inadequate for the thin tailed standardized error distributions and, consequently, the empirical approach fails to dynamically predict the BP exchange rate risk accurately at 95% confidence on long U.S. dollar positions.

A natural question that arises from the above discussion is that the dynamic forecasting problem exhibited by the BP is not observed in the remaining currencies, although, the estimated t-distributions are fatter at the tails relative to the distributions of the standardized residuals across the currencies (the estimated t-distribution forecasts fail the distribution tests for all the currencies because they are too cautious). Some explanations of this can be obtained by reference to the densities at the tails of the forecast error distributions and the amount of time variations in the conditional variances. The results from Table 3.14 show that, with the exception of the CD, the selected t-distribution likelihoods produce the least time variation in the forecast variances of all the exchange rates. The results make sense relative to the normal distribution likelihoods but not the estimated t-distribution likelihoods. This is because maximization of the likelihood function ensures that, within the constraints of the GARCH model, the resulting empirical distribution of the standardized errors represents the assumed distribution as much as possible. Thus, the normal distribution will be expected to produce the least densities at the tails of its standardized residuals and, hence, would require more time variation in the conditional variances to model a given empirical leptokurtic distribution. The results from the normal and the selected t-distributions confirm this, as the conditional variances from the normal estimates produce higher time variations across the currencies. The estimated t-distribution likelihoods, however, do not generate the expected results (except for the CD). The reason for this can be unearthed by reference to how the two t-distributions were "chosen" to describe the distribution of the errors. The selected t-distributions were chosen to represent the 1st and the 99th percentile empirical values as much as possible without regard to the distribution of the points in between. The goal here, therefore, is to model the empirical tail fatness and not the center or the shoulders of the empirical distribution. The estimated t-distribution, on the other hand, were estimated to fit the entire distribution of the errors as much as possible. Since the selected t-distributions, even without time-varying conditional variances, have enough densities at the tails to model the exchange rates at 98% central confidence, the fatter tailed estimated t-distributions cannot be justified on grounds of modeling the tails. Evidence in support of this argument is that the more than expected time variations in the conditional variances generated by the estimated t-distributions, to the contrary, result in less mass at the tails of the residuals. Table 3.15 reports the average empirical standardized VaR forecast values at 95% and 99% confidence for both short and long positions. The empirical interval forecasts generated by the estimated t-distribution likelihoods are the smallest across the exchange rates. This suggests that relatively high time variations in the conditional variances will be necessary for the empirical distributions obtained with the estimated t-distribution likelihoods to properly forecast VaR. The high time variations generated by the estimated t-distributions in the conditional variances are, therefore, sufficient to dynamically model the VaRs of the exchange rates accurately except the BP.

The estimated t-distributions produce the highest time variations in the conditional variances of the BP, the DM, and the FF, while the normal generate the highest time variations for the CD and the JY. For the CD, the estimated t-distribution produces the least time variation in the forecast variances as expected. This, couple with the smallest width of the empirical forecast intervals, suggests that the VaR forecasts of the CD exchange rate may not be accurate under the estimated t-distribution likelihoods. The results confirm this. The empirical forecasts produced by the estimated t-distribution likelihood under forecast the risk of the CD the most. It under predicts the CD exchange rate risk by 16% on long U.S. dollar positions and 24% on shorts.

3.6. DISCUSSION

3.6.1. The Theoretical Forecasts

We use three theoretical forecasts in this study: the normal, the selected t-distribution and the estimated t-distribution. Among these distributions, the normal distributions are the most confident and the estimated t-distributions are the most cautious. It is, therefore, not surprising that the normal VaR forecasts yield the highest exceedence ratios and the estimated t-distributions produce the least exceedence ratios across the currencies. The performances of the different theoretical distributions differ at the different levels of confidence. At the 95% level, the normal distribution forecasts are not over confident. This is evidenced by the normal VaR forecasts overestimating the risks of some of the exchange rates (both positions for the JY and longs for the BP). Indeed, at the 95% confidence level, the normal distributions' values at risk outperform their t-distributions' counterparts. This is consistent with the empirical evidence that, at this level of confidence, the normal distribution is sufficiently fat to properly model the exchange rates' risks. Since the t-distributions are fatter relative to the normal distribution at the tails, the numbers of standard deviations associated with the 95% confidence level are higher than their empirical analogs. Thus, both t-distributions were too cautious leading to over-estimation of the risks of the exchange rates at the 95% confidence level. The most cautious distribution (the estimated t-distribution) produces the worst VaR forecasts (its forecasts failed all the unconditional likelihood ratio tests), and its best forecasts over estimate the risk of the GM exchange rate by 43% on long U.S. dollar positions.[49]

The numbers of standard deviations corresponding to the 99% confidence level of the selected t-distributions best approximate their standardized residual counterparts, as the t-distributions were selected so as to match the 1st and the 99th percentile standardized empirical values as much as possible. Consequently, the selected t-distributions produced the best forecasts among the three theoretical forecasts at the 99% confidence level. At this level, the normal distributions are too confident (they failed all the available unconditional likelihood ratio tests) and the estimated t-distributions are too cautious (they failed three of the possible ten unconditional likelihood ratio tests). It is clear from these results that the performances of the models depend critically on the relationship between the theoretical number of standard deviations at the associated confidence level and its standardized residual analog. If the theoretical number of standard deviations is higher (theoretical distribution is fat tailed relative to the error distribution), the model tends to over predict the risks of the exchange rates and, if the theoretical number of

[49] Christofferssen (1998) also find that the t-distribution/GARCH failed to predict VaR accurately because it was too cautious.

standard deviations is lower, the model tends to under predict the risks.[50] It does not, therefore, appear that poor VaR forecasts crucially depend on the traditional measure of kurtosis, but rather, on kurtosis as measured by the number of standard deviations at the given confidence level. For instance, the empirical kurtosis of all the exchange rates are in excess of the normal value, but the normal distribution yields the best forecasts at 95% level of confidence among the theoretical distributions.[51] This finding, therefore, support Duffie and Pan (1997) contention that, for the VaRs of direct risk exposures, a better measure of kurtosis is the number of standard deviations associated with the confidence level. Our results together with Duffie and Pan's contention suggest that, for the purposes of forecasting VaRs, research should be focused on modeling excess kurtosis as measured by the number of standard deviations associated with the particular level of confidence, and not the standard measure of excess kurtosis. Given this, a VaR forecast which uses a theoretical distribution that is estimated such that the number of standard deviations of the theoretical distribution and the sample residuals are as equal as possible at the relevant confidence level is likely to be very successful in forecasting VaR.

The study also sheds light on the effects of imposing different theoretical distributions on the amount of kurtosis that is captured by time-varying variances. An imposition of a particular theoretical distribution on the GARCH model ensures that the parameters of the model are estimated such that the resulting standardized errors resemble the assumed theoretical distribution as much as possible. This suggests that the normal distribution would generate standardized errors that are least dense at the tails and

[50] For example, the normal distribution under predicts the risk of the exchange rates at 99%, but not at 95%, while the selected t-distribution over predicts the risks of the exchange rates at 95% confidence but not at 99%.

the estimated t-distributions will produce the fattest tailed standardized residuals. For any given empirical kurtosis, the higher the density at the tails of the standardized residuals, the less would be the time variation require in the conditional variances to model the empirical kurtosis. This argument suggests that the normal likelihoods would generate the highest time variations in the conditional variances and the estimated t-distributions would produce the least time variations across the currencies. The result from Table 3.14 does not support this contention. Indeed, the estimated t-distributions produce the highest time variations in the forecast variances for the BP, DM, and the FF. On the other hand, as expected, the normal distributions produce higher time variations than the selected t-distribution across the exchange rates. The reason for this is that while the selected t-distributions were chosen to model the empirical 1st and 99th percentile values (tails), the estimated t-distributions were obtained by maximizing the likelihood function with respect to the entire sample. Thus, the results from the selected t-distributions, which focus on modeling the tails, are consistent with our expectation while those from the estimated t-distributions are not because of the influence of modeling of the center and the shoulders of the empirical distributions. The implication of this is that tail fatness of the samples alone is not sufficient to elicit systematic changes in the variations of the conditional variances. This observation suggests that to the extent that the data points in the center and the shoulders of the empirical distribution do not significantly influence the estimations, an imposition of a fatter tail theoretical distribution on the errors would result in less time variation in the conditional variances, as the normal and the selected t-distributions demonstrate.

[51] Also, the selected t-distributions over predict the risks of the CD, the FF, and the JY when the sample statistics show that there are some unaccounted for kurtoses in these exchange rates.

3.6.2. The Empirical Forecasts

The parameter estimates are different for the three distributions because they are estimated such that, for any given sample, the standardized residuals fit the assumed distribution as much as possible. Thus, the different distributional assumptions produced different parameter estimates of the GARCH model. The different estimates of the variance equation parameters would produce different time variations in the predicted variances. These different time variations produce different dynamic VaR forecasts by contracting and spreading the assumed distribution to different degrees for given shocks to the exchange rates. The higher the time variation in the conditional variances, the more the assumed distribution spreads during volatile periods and, also, the more it contracts during tranquil periods. Thus, the empirical forecasts generated by the different likelihoods are driven by different dynamic variance forecasts produce by the different likelihood functions. Since the empirical distributions by construction account for all the empirical kurtosis, the differences in the VaR forecasts resulting from the use of the different likelihoods are only due to the apportionment of the empirical kurtosis between time varying conditional variances and the forecast empirical distribution. *Ceteris paribus*, the more the kurtosis explained by the time variation of the conditional variances the less the density of the empirical distribution at the tails. This assertion is confirmed by the time variation in the conditional variances and the resulting empirical bounds of the normal and the selected t-distributions.[52] For the BP, the estimated t-distribution produces the highest time variation in the conditional variances and the least densities at the tail of the empirical distributions (see Table 3.15). The results show that the BP exchange rate fails the independence test on short position at 10% when the estimated t-distribution likelihood is used but passes the test when the normal and the selected t-distribution likelihoods are used. Thus, the results suggest that the excess kurtosis of the BP exchange rate is mainly due to fat-tailed distribution rather than changing volatilities.[53] For the other exchange rates, it appears the source of the empirical kurtosis is not crucial in determining the accuracy of the VaR forecasts (similar empirical forecasts are obtained with the different likelihoods).

Our results also show that the use of the empirical methodology can pose problems if the dynamic variance equation, which is the best estimate of the relationship among the conditional variances given the assumed theoretical distribution, is combined with empirical distributions that are different from the theoretical one underlying the variance equation. Forecasting problems using the empirical distribution can arise when the densities at the tails are inappropriate or when the *ad hoc* combination of the estimated variance equations and the empirical distributions is unsuitable for the purposes of forecasting VaR. The former problem is depicted by the results of the JY exchange rate. These results show disparity in the theoretical distributions' forecasts and their empirical counterparts at the 95% confidence level on short U.S. dollar positions. While the theoretical distributions over predict the risk, the empirical distributions under predict it. The empirical 90% central confidence bands are less than their theoretical counterparts and, therefore, explain why the theoretical distributions over predict the JY exchange rate

---

[52] As pointed out earlier, the results from the estimated t-distributions are noisy because of modeling the center and the shoulders of the distributions and, hence, they do not necessarily conform to the expected outtcome.

risk. The likelihood ratio tests shed some light on the sources of the under predictions resulting from the empirical distributions' forecasts. The results suggest that the under predictions of the JY risk on short U.S. dollar positions are due to thin densities at the tails (the unconditional likelihood ratio statistics are all higher than 2). Although the likelihood ratio statistics are higher than 2, only the forecasts under the estimated t-distribution likelihood fails the test at 10% level of significance (it can be seen from Table 3.15 that, for the JY, the estimated t-distribution produces empirical distribution forecasts with the least densities at the tails). An argument against this kind of problem is that the theoretical distributions are fatter at the tails than the empirical distribution of the JY and, hence, it is not necessary to employ the empirical distribution procedure to model tail fatness.

The latter problem is demonstrated by the BP exchange rate. The empirical VaR forecasts of the BP pass the dynamic variance (independence) tests when the normal and the selected t-distribution likelihoods were used, but not when the t-distribution likelihood is estimated at the 95% confidence level. The reason for this is that the estimated t-distribution is most different (compare to the normal and the selected t-distribution) from the distribution of its errors and, therefore, produces the highest mismatching of forecast empirical distributions and estimated variance equations. Thus, whereas the estimates of the variance equation together with the estimated t-distribution predictions produce VaR forecasts that pass the independence test, the same estimates of the variance equation are unsuitable for the empirical distributions, and this results in the failure of the independence test at 5% level of significance.

[53] Brooks (1996), using daily data on the BP from January 2, 1974 to July 1, 1994, finds that only very modest improvements in the forecasts can be achieved by using the best linear and non-linear univariate

3.6.3. The Theoretical versus Empirical Forecasts

The empirical forecasts account for all the empirical kurtosis irrespective of the likelihood function used and, therefore, are more likely to produce better VaR estimates. From the results, it is clear that the empirical forecasts dominate the theoretical ones when the theoretical distributions do not appropriately capture the empirical kurtosis as measured by the number of standard deviations associated with the particular confidence level. Thus, the empirical approach does not clearly dominate the GARCH/normal and the GARCH/selected-t at 95% and 99% confidence levels, respectively. The implication of this is that a superior empirical VaR performance is a signal that the distributional assumption underlying the GARCH model is incorrect, at least, at the relevant confidence level.[54] This raises the question of why a wrong distributional assumption would be used to estimate the parameters of the GARCH equations in the first place. An answer to this is that the parameter estimates of the GARCH model are still consistent under some weak conditions even if the distributional assumption is inappropriate (Quasi Maximum Likelihood Estimation).[55] Engle and González-Rivera (1991) have shown that the efficiency loss can be high if the true distribution is very different from the assumed distribution underlying the parameter estimates (for example, if the true distribution is a t-distribution with 5 degrees of freedom, then an imposition of a normal would result in about 58% loss in efficiency). This loss in efficiency will reduce the forecasting ability of

---

models rather than the random walk model.

[54] Thus, the empirical approach can be viewed as an *ad hoc* way of dealing with a poor distributional assumption.

[55] Bollerslev and Woodbridge (1988) show that under a correct specification of the first and second moments, consistent estimates of the parameters of the model can be obtained by maximizing a normal likelihood function even if the true density is not normal.

the models.[56] Consequently, they propose using non-parametric methods to estimate the distribution of the errors rather than imposing some theoretical distribution. Their technique is likely to improve the empirical forecasting methodology since it avoids combining parametric and non-parametric methods in an *ad hoc* fashion.

3.7. CONCLUSION

The study reveals certain important results regarding the VaR forecasts of the five exchange rates examined. First, the study shows that the traditional measure of kurtosis may not be appropriate for VaRs of direct exposures. Traditionally measured, all the exchange rates have excess kurtosis relative to the normal, but we find that the normal distribution produced VaR forecasts that are superior to the two fatter tailed t-distributions at 95% confidence level for all the exchange rates. The explanation for this is that relative to the number of standard deviations at the 95% confidence, the empirical distributions of the exchange rates are not leptokurtic relative to the normal benchmark. This finding supports the contention of Duffie and Pan (1997) that, for the purposes of forecasting VaRs of direct exposures, a more pertinent measure of kurtosis should be the number of standard deviations represented by the associated confidence level. At the 99% level of confidence, however, the normal distribution's number of standard deviations is less than the empirical 99^{th} percentile values of the standardized errors for both positions across the currencies. This suggests that the exchange rates exhibit excess kurtosis relative to the normal distribution at this confidence level. It is, therefore, not surprising that the normal distribution was too confident at 99% VaR confidence level and, as a

[56] Since the GARCH methodology requires prediction of the future volatilities, potential problems can arise if the wrong distributional assumption is used in estimating the GARCH parameters.

result, underestimates the risks of the exchange rates. These results suggest that modeling kurtosis as measured by the number of standard deviations associated with the confidence level would improve VaR forecasts while modeling kurtosis as is traditionally measured may not produce improvements.

In view of this finding, VaR estimates based on distributions that are selected such that the number of standard deviations at the relevant confidence level matches the counterpart of the standardized residual are likely to generate good VaR forecasts. In this study, we find that the selected t-distribution yielded VaR forecasts that are superior to the normal and the estimated t-distribution forecasts at the 99% confidence level. This is because the selected t-distributions were chosen such that their theoretical values are as close as possible to the standardized empirical 98% central confidence values. Thus, better results would be obtained if they were selected to match a particular position (long or short) as this would minimize the errors in the VaR due to skewness. A more general way of dealing with skewness is to use models, such as the EGARCH, which explicitly models skewness. It appears in this study that skewness does not have important bearing on the results since we do not find systematic differences between the results from the empirical approach, which controls for skewness, and the results from the theoretical distributions (normal and t-distribution).

The study also demonstrates that, in addition to proper modeling of the empirical kurtosis, its division between the imposed distribution and time-varying variances is important in accurately forecasting some risk factors. With the BP exchange rate, for instance, the results indicate that it can be properly modeled by appropriate fat-tail distributions rather than more general volatility models. This finding is important because it suggests that progress can be made in accurately forecasting the BP exchange rate risk by focusing research on identifying appropriate fat-tail distributions. For the other exchange rates, the results do not suggest that the distribution of the unconditional kurtosis between time-varying volatility and the assumed distribution is important in forecasting VaR.

We also find that the empirical distribution methodology does not always dominate the theoretical forecasts. In particular, where the theoretical number of standard deviations associated with the level of VaR confidence is close to the sample analog, both approaches perform well and there is no clear dominance. For example, we find that the empirical forecasts are not superior to the normal forecasts at 95% confidence level, and they are also not superior to the selected t-distribution forecasts at 99%. This indicates that the empirical forecasting methodology dominates when the distributional assumption is incorrect at the relevant confidence level. This was the case with the estimated t-distributions which are clearly over cautious at the two confidence levels used in this study. However, under these circumstances, the empirical forecast could also potentially produce poor VaR forecasts, which would be captured by clustering of VaR violations, if the estimated dynamics of the variance equation is not consistent with the empirical distribution as evidenced by the BP exchange rate. In these situations, estimating the parameters using non-parametric distributions along the lines suggested by Engle and González-Rivera (1991) would counter this potential problem.

Table 3.1. Summary Statistics of Daily Exchange Rate Returns (In-Sample)

Statistics	BP	DM	CD	FF	JY
Mean	.0023	-.0066	.0147	-.0060	-.0281
Standard deviation	.6865	.7386	.2700	.7145	.6618
Skewness	.2673	.1112	.2021	-.0846	-.4015
Kurtosis	5.276	4.910	5.101	5.231	7.944
Maximum	3.31%	3.46%	1.25%	3.46%	4.13%
Minimum	-2.82%	-3.50%	-1.30%	-4.83%	-4.95%
		Test for Normality			
$\chi^2(57)$	313*	198*	334*	354*	215*
Jarque-Bera LM [($\chi^2(2)$)	293*	198*	246*	563*	1350*
		Test of equal yearly distribution of returns			
Kruskal-Wallis Test	712*	831*	1057*	1142*	1018*

Sample Size is 1303. The standard deviation of skewness is .0678 and that of kurtosis is .1355. At 1%, $\chi^2(60) = 88.4$ and $\chi^2(2) = 9.21$. (* indicates significant at 1%).

Table 3.2. Summary Statistics of Daily Exchange Rate Returns (Out-of-Sample)

Statistics	BP	DM	CD	FF	JY
Mean	-.0023	.0173	.0022	.0153	.0019
Standard deviation	.4511	.6055	.3223	.5884	.8459
Skewness	-.0808	-.1341	-.1448	-.4140	-1.1562
Kurtosis	5.474	5.469	5.476	7.121	12.095
Maximum	2.37%	3.13%	1.40%	3.12%	3.95%
Minimum	-2.24%	-2.80%	-1.57%	-4.14%	-7.68%
		Test for Normality			
$\chi^2(57)$	432*	194*	216*	208*	326*
Jarque-Bera LM [($\chi^2(2)$)	330*	331*	334*	951*	4748*
		Test of equal yearly distribution of returns			
Kruskal-Wallis Test	712*	831*	1057*	1142*	1018*

Sample Size is 1660. The standard deviation of skewness is .0678 and that of kurtosis is .1355. At 1%, $\chi^2(60) = 88.4$ and $\chi^2(2) = 9.21$. (* indicates significant at 1%).

Table 3.3. Autocorrelations and Daily Exchange Rate Returns and Squared Returns (In-Sample)

	Autocorrelation of Returns					Autocorrelation of Squared Returns				
Lag	BP	DM	CD	FF	JY	BP	DM	CD	FF	JY
1	0.098	0.017	0.051	0.021	-0.013	0.146	0.061	0.058	0.075	0.126
2	0.006	-0.018	-0.023	-0.014	-0.053	0.138	0.094	0.034	0.057	0.04
3	0.004	-0.020	-0.023	-0.012	-0.037	0.141	0.062	0.057	0.048	0.065
4	0.056	0.022	0.002	0.032	0.006	0.071	0.060	0.060	0.051	0.073
5	0.063	0.019	0.038	0.005	0.003	0.112	0.028	0.097	0.027	0.099
6	-0.035	-0.076	-0.031	-0.075	-0.085	0.093	0.090	0.075	0.058	0.096
7	-0.053	0.016	-0.004	0.023	0.054	0.071	0.032	0.052	0.026	0.107
8	0.036	-0.017	-0.010	-0.030	-0.009	0.040	0.052	0.070	0.039	0.044
9	0.031	0.006	-0.016	-0.003	0.028	0.120	0.013	0.085	0.025	0.049
10	0.048	0.050	-0.003	0.032	0.063	0.076	0.078	0.017	0.065	0.124
11	-0.004	-0.017	0.032	0.012	-0.012	0.184	0.091	0.031	0.104	0.121
12	0.011	0.021	0.032	0.022	0.028	0.100	0.026	-0.025	0.006	0.021
	Adjusted Box-Pierce					Adjusted Box-Pierce				
Q(40)	66.55	35.27	34.48	37.33	45.39	369	136	119	155	44.55
	(.005)	(.687)	(.716)	(.591)	(.257)	(.000)	(.000)	(.000)	(.000)	(.286)

Asymptotic standard error for the autocorrelation estimates is .028. Adjusted Box-Pierce significant probabilities in parentheses.

Table 3.4. GARCH Estimates (Normal Likelihood)

Statistics	BP	DM	CD	FF	JY
Mean Equation					
Constant	-.0274	-.0159	.0071	-.0165	-.0288
	(.0162)	(.0191)	(.0068)	(.0188)	(.01841)
Variance Equation					
Constant	.0058	.0123	.0010	.0080	.0094
	(.0024)	(.0058)	(.0004)	(.0048)	(.0047)
α	.0602	.0376	.0346	.0322	.0222
	(.0114)	(.0091)	(.0090)	(.0093)	(.0084)
β	.9281	.9391	.9522	.9519	.9564
	(.0140)	(.0174)	(.0126)	(.0169)	(.0167)
ML	-1268.78	-1418.82	-103.68	-1371.39	-1292.81

Sample Size is 1303. Standard errors in parentheses.

Table 3.5. Summary Statistics of Standardized Residuals (GARCH/normal)

Statistics	BP	GM	CD	FF	JY
Mean	.0254	.0049	.0120	-.0029	-.0079
Skewness	.0475	.0986	.3078	-.2362	-.4987
Kurtosis	4.441	4.417	5.833	6.1011	7.9845
Maximum	3.9514	4.5448	4.648	4.0101	5.8484
Minimum	-4.2465	-4.0219	-6.108	-7.2019	-7.8027
5^{th} percentile	-1.63	-1.70	-1.50	-1.65	-1.62
95^{th} percentile	1.65	1.65	1.66	1.63	1.46
1^{st} percentile	-2.69	-2.72	-2.56	-2.64	-2.97
99^{th} percentile	2.45	2.54	2.65	2.52	2.21
		Test for Normality			
$\chi^2(57)$	171	154	270	235	203
Jarque-Bera LM [$(\chi^2(2))$	112	110	451	528	1389

Sample Size is 1303. The standard deviation of skewness is .0678 and that of kurtosis is .1355. At 1%, $\chi^2(60) = 88.4$ and $\chi^2(2) = 9.21$.

Table 3.6. Various Statistics of the Forecasts of the Normal Distribution Likelihood (Empirical and Theoretical Forecasts)

Forecast Dist.		Confidence Level	Observed Exceedence	Expected Exceedence	Exceedence Ratio	Degree of Exceedence	MAD	RMSE
BP	Normal	95%(long)	81	83	0.98	0.15351	0.00507	0.00579
		95%(short)	83	83	1.00	0.12496	0.00495	0.00572
		99%(long)	27 (4)	16.6	1.63	0.05366	0.00698	0.00765
		99%(short)	18 (2)	16.6	1.08	0.03338	0.00688	0.00758
	Empirical	95%(long)	86	83	1.04	0.16587	0.00496	0.00569
		95%(short)	85	83	1.02	0.13086	0.00492	0.00570
		99%(long)	18 (1)	16.6	1.08	0.03957	0.00788	0.00852
		99%(short)	14 (2)	16.6	0.84	0.02641	0.00737	0.00807
GM	Normal	95%(long)	89	83	1.07	0.68013	0.01936	0.02226
		95%(short)	91	83	1.10	0.51987	0.01834	0.02166
		99%(long)	34 (9)	16.6	2.03	0.22269	0.02644	0.02923
		99%(short)	25 (3)	16.6	1.51	0.13394	0.02550	0.02857
	Empirical	95%(long)	86	83	1.04	0.65913	0.01953	0.02246
		95%(short)	90	83	1.08	0.52908	0.01832	0.02162
		99%(long)	16 (3)	16.6	0.96	0.10549	0.03095	0.03370
		99%(short)	15 (2)	16.6	0.90	0.08330	0.02805	0.03099
CD	Normal	95%(long)	85	83	1.02	0.23067	0.00780	0.00911
		95%(short)	93	83	1.12	0.23531	0.00763	0.00903
		99%(long)	27 (4)	16.6	1.63	0.07082	0.01074	0.01199
		99%(short)	25 (2)	16.6	1.51	0.05851	0.01054	0.01191
	Empirical	95%(long)	90	83	1.08	0.24793	0.00763	0.00899
		95%(short)	95	83	1.14	0.24251	0.00755	0.00897
		99%(long)	17 (3)	16.6	1.02	0.03844	0.01213	0.01341
		99%(short)	13 (1)	16.6	0.78	0.03324	0.01191	0.01324
FF	Normal	95%(long)	88	83	1.06	2.26566	0.06489	0.07497
		95%(short)	83	83	1.00	1.57867	0.06145	0.07299
		99%(long)	33 (7)	16.6	1.99	0.77688	0.08870	0.09861
		99%(short)	23 (2)	16.6	1.38	0.40925	0.08565	0.09647
	Empirical	95%(long)	89	83	1.07	2.32432	0.06403	0.07454
		95%(short)	88	83	1.06	1.58536	0.06129	0.07266
		99%(long)	16 (2)	16.6	0.96	0.46529	0.09919	0.10923
		99%(short)	17 (3)	16.6	1.02	0.26263	0.09336	0.10364
JY	Normal	95%(long)	73	83	0.88	64.9444	1.51555	1.76915
		95%(short)	74	83	0.89	33.8062	1.44946	1.76586
		99%(long)	34 (1)	16.6	2.03	35.6967	2.07906	2.33040
		99%(short)	19 (4)	16.6	1.14	11.7490	2.02162	2.32686
	Empirical	95%(long)	74	83	0.89	66.7612	1.49431	1.75115
		95%(short)	103	83	1.24	48.5015	1.29333	1.61134
		99%(long)	20 (0)	16.6	1.20	20.9291	2.72534	2.99279
		99%(short)	22 (4)	16.6	1.32	14.2973	1.96342	2.25618

Observed exceedences in the last 250 days are in brackets. MAD is mean absolute deviation.

Table 3.7. Likelihood Ratio Test Statistics of the Normal Forecasts

Forecast Dist		Normal			Empirical		
Test		Unconditional	Independence	Conditional	Unconditional	Independence	Conditional
Lower Tail	**BP**						
	95%	0.022	3.244*	3.266	0.049	2.550	2.599
	99%	2.403	0.238	2.641	0.050	0.752	0.803
	GM						
	95%	0.194	0.143	0.337	0.049	0.236	0.285
	99%	**6.140**	0.746	**6.886**	0010	N/A	N/A
	CD						
	95%	0.022	1.211	1.233	0.263	2.066	2.329
	99%	2.403	N/A	N/A	0.004	N/A	N/A
	FF						
	95%	0.135	0.171	0.306	0.194	0.143	0.337
	99%	**5.521**	0.819	**6.340**	0.070	N/A	N/A
	JY						
	95%	0.573	0.937	1.510	0.462	0.864	1.326
	99%	**6.140**	1.948	**8.088**	0.287	0.603	0.889
Upper Tail	**BP**						
	95%	0.0	0.079	0.079	0.022	0.014	0.036
	99%	0.050	0.752	0.803	0.189	N/A	N/A
	GM						
	95%	0.342	0.101	0.444	0.263	0.081	0.344
	99%	1.614	N/A	N/A	0.070	N/A	N/A
	CD						
	95%	0.531	0.274	0.805	0.759	0.524	1.283
	99%	1.614	1.587	3.201	0.370	3.781*	4.151
	FF						
	95%	0.0	0.168	0.168	0.135	0.048	0.183
	99%	0.966	N/A	N/A	0.004	N/A	N/A
	JY						
	95%	0.462	0.280	0.742	2.053	1.080	3.133
	99%	0.145	N/A	N/A	0.699	N/A	N/A

At 5%, $\chi^2(1) = 3.84$ $\chi^2(2) = 5.99$. At 10% $\chi^2(1) = 2.71$ $\chi^2(2) = 4.61$ The highlighted likelihood ratio statistics are significant at 5% and those with * are significant at 10%.

Table 3.8. Selected and Estimated t-distributions GARCH Estimates

Currency	BP	DM	CD	FF	JY
Selected df	**14**	**13**	**11**	**14**	**10**
Mean Eq. Constant	-.0296 (.0162)	-.0179 (.0191)	-.0004 (.0064)	-.0180 (.0177)	-.0167 (.0163)
Variance Eq. Constant	.0066 (.0030)	.0102 (.0052)	.0003 (.0003)	.0087 (.0050)	.0054 (.0032)
α	.0661 (.0142)	.0344 (.0089)	.0366 (.0090)	.0347 (.0102)	.0195 (.0067)
β	.9170 (.0180)	.9436 (.0160)	.9560 (.0110)	.9442 (.0185)	.9639 (.0130)
ML	-497.32	-650.83	688.86	-591.18	-482.31
Estimated df	**5.06**	**5.36**	**4.33**	**4.94**	**4.08**
Mean Eq. Constant	-.0306 (.0154)	-.0186 (.0189)	-.0026 (.0057)	-.0207 (.0190)	-.0082 (.0161)
Variance Eq. Constant	.0077 (.0039)	.0094 (.0054)	.0002 (.0002)	.0086 (.0055)	.0056 (.0040)
α	.0813 (.0175)	.0358 (.0103)	.0410 (.0106)	.0381 (.0125)	.0259 (.0092)
β	.9083 (.0192)	.9481 (.0158)	.9603 (.0101)	.9465 (.0193)	.9630 (.0134)
ML	-489.82	-642.51	707.64	-577.64	-465.54

Sample Size is 1303. Standard errors in parentheses.

Table 3.9. Summary Statistics of the Standardized Residuals from the t-distributions estimates (In-Sample)

Currency	BP	DM	CD	FF	JY
Selected df	**14**	**13**	**11**	**14**	**10**
Mean	.0292	.0076	.0469	.0053	-.0291
Skewness	.0330	.0969	.3786	-.2465	-.5217
Kurtosis	4.47 (3.86)	4.42 (3.67)	6.47 (3.86)	6.17 (3.27)	8.06 (4)
Maximum	4.0857	4.6558	5.5981	4.1986	6.0894
Minimum	-4.4208	-4.1462	-6.6599	-7.5492	-8.4143
5^{th} percentile	-1.66 (-1.80)	-1.74 (-1.77)	-1.54 (-1.80)	-1.71 (-1.76)	-1.76 (-1.81)
95^{th} percentile	1.68 (1.80)	1.69 (1.77)	1.73 (1.80)	1.70 (1.76)	1.53 (1.81)
1^{st} percentile	-2.71 (-2.62)	-2.77 (-2.65)	-2.61 (-2.72)	-2.69 (-2.62)	-3.21 (-2.76)
99^{th} percentile	2.54 (2.62)	2.61 (2.65)	2.86 (2.72)	2.61 (2.62)	2.37 (2.76)
Estimated df	**5.06**	**5.36**	**4.33**	**4.94**	**4.08**
Mean	.0288	.0079	.0560	.0084	.0079
Skewness	.0193	.0919	.4366	-.2500	.0919
Kurtosis	4.50 (8.66)	4.43 (7.41)	6.79 (21.2)	6.16 (9.38)	4.43 (78)
Maximum	3.9618	4.4747	5.6487	3.9797	4.4747
Minimum	-4.2913	-4.0256	-6.1928	-7.1778	-4.0256
5^{th} percentile	-1.58 (-2.01)	-1.67 (-2.01)	-1.41 (-2.13)	-1.63 (-2.01)	-1.67 (-2.13)
95^{th} percentile	1.62 (2.01)	1.62 (2.01)	1.61 (2.13)	1.63 (2.01)	1.62 (2.13)
1^{st} percentile	-2.57 (-3.36)	-2.67 (-3.36)	-2.40 (-3.75)	-2.59 (-3.36)	-2.67 (-3.75)
99^{th} percentile	2.45 (3.36)	2.52 (3.36)	2.71 (3.75)	2.48 (3.36)	2.52 (3.75)

The theoretical values are reported in brackets.

Table 3.10. Various Statistics of the Forecasts of the Selected t-distribution Likelihood (Empirical and Theoretical Forecasts)

Forecast Dist.		Confidence Level	Observed Exceedence	Expected Exceedence	Exceedence Ratio	Degree of Exceedence	MAD	RMSE
BP	t-distrib	95%(long)	74	83	0.89	0.13921	0.00524	0.00597
		95%(short)	79	83	0.95	0.10995	0.00512	0.00590
		99%(long)	19 (4)	16.6	1.14	0.04003	0.00758	0.00825
		99%(short)	14 (1)	16.6	0.84	0.02488	0.00748	0.00817
	Empirical	95%(long)	86	83	1.04	0.16671	0.00494	0.00569
		95%(short)	87	83	1.05	0. 12794	0.00495	0.00573
		99%(long)	19 (3)	16.6	1.14	0.03846	0.00783	0.00851
		99%(short)	16 (2)	16.6	0.96	0.02715	0.00736	0.00806
GM	t-distrib	95%(long)	78	83	0.94	0.58973	0.02019	0.02319
		95%(short)	79	83	0.95	0.44873	0.01919	0.02258
		99%(long)	22 (3)	16.6	1.33	0.13489	0.02910	0.03203
		99%(short)	17 (2)	16.6	1.02	0.07888	0.02820	0.03137
	Empirical	95%(long)	85	83	1.02	0.64457	0.01964	0.02267
		95%(short)	91	83	1.10	0.53352	0.01834	0.02172
		99%(long)	16 (2)	16.6	0.96	0.09683	0.03116	0.03403
		99%(short)	16 (2)	16.6	0.96	0.07566	0.02828	0.03136
CD	t-distrib	95%(long)	70	83	0.84	0.19102	0.00824	0.00957
		95%(short)	81	83	0.98	0.19216	0.00806	0.00950
		99%(long)	17 (3)	16.6	1.02	0.03603	0.01212	0.01342
		99%(short)	15 (1)	16.6	0.90	0.03189	0.01194	0.01333
	Empirical	95%(long)	93	83	1.12	0.24494	0.00768	0.00907
		95%(short)	97	83	1.17	0.24626	0.00756	0.00900
		99%(long)	18 (3)	16.6	1.08	0.04081	0.01209	0.01343
		99%(short)	18 (1)	16.6	1.08	0.03286	0.01194	0.01333
FF	t-distrib	95%(long)	81	83	0.98	2.05721	0.06656	0.07714
		95%(short)	77	83	0.93	1.44287	0.06321	0.07524
		99%(long)	19 (3)	16.6	1.14	0.54760	0.09551	0.10615
		99%(short)	20 (2)	16.6	1.20	0.26754	0.09257	0.10411
	Empirical	95%(long)	87	83	1.05	2.22589	0.06455	0.07536
		95%(short)	82	83	0.99	1.60031	0.06149	0.07317
		99%(long)	18 (2)	16.6	1.08	0.51355	0.09925	0.10979
		99%(short)	16 (2)	16.6	0.96	0.24690	0.09400	0.10485
JY	t-distrib	95%(long)	68	83	0.82	63.8434	1.53509	1.77942
		95%(short)	66	83	0.80	31.5350	1.46758	1.77579
		99%(long)	30 (1)	16.6	1.81	29.4099	2.25223	2.49228
		99%(short)	13 (2)	16.6	0.78	8.69370	2.19869	2.48832
	Empirical	95%(long)	80	83	0.96	67.6262	1.49220	1.74261
		95%(short)	107	83	1.29	48.4500	1.28639	1.60053
		99%(long)	19 (1)	16.6	1.14	19.4985	2.70459	2.95974
		99%(short)	21 (4)	16.6	1.27	13.1401	1.92709	2.21294

Observed exceedences in the last 250 days are in brackets. MAD is mean absolute deviation.

Table 3.11. Various Statistics of the Forecasts of the Estimated t-distribution Likelihood (Empirical and Theoretical Forecasts)

Forecast Dist.		Confidence Level	Observed Exceedence	Expected Exceedence	Exceedence Ratio	Degree of Exceedence	MAD	RMSE
BP	t-distrib	95%(long)	46	83	0.55	0.08422	0.00620	0.00691
		95%(short)	37	83	0.45	0.05518	0.00609	0.00684
		99%(long)	5 (1)	16.6	0.30	0.01683	0.01018	0.01082
		99%(short)	2 (0)	16.6	0.12	0.00549	0.01009	0.01074
	Empirical	95%(long)	88	83	1.06	0.16420	0.00496	0.00571
		95%(short)	88	83	1.06	0.12604	0.00495	0.00574
		99%(long)	17 (2)	16.6	1.02	0.03769	0.00787	0.00854
		99%(short)	14 (2)	16.6	0.84	0.02553	0.00737	0.00809
GM	t-distrib	95%(long)	47	83	0.57	0.31672	0.02415	0.02714
		95%(short)	41	83	0.49	0.21153	0.02319	0.02651
		99%(long)	4 (0)	16.6	0.24	0.02064	0.03943	0.04235
		99%(short)	3 (0)	16.6	0.18	0.01450	0.03860	0.04167
	Empirical	95%(long)	87	83	1.05	0.63322	0.01975	0.02286
		95%(short)	86	83	1.04	0.51648	0.01841	0.02183
		99%(long)	17 (3)	16.6	1.02	0.10226	0.03127	0.03425
		99%(short)	17 (2)	16.6	1.02	0.08067	0.02839	0.03155
CD	t-distrib	95%(long)	39	83	0.47	0.09398	0.01005	0.01136
		95%(short)	39	83	0.47	0.08030	0.00985	0.01128
		99%(long)	2 (0)	16.6	0.12	0.00474	0.01741	0.01871
		99%(short)	0 (0)	16.6	0	0	0.01722	0.01863
	Empirical	95%(long)	96	83	1.16	0.24440	0.00769	0.00909
		95%(short)	103	83	1.24	0.24355	0.00756	0.00902
		99%(long)	19 (1)	16.6	1.14	0.04332	0.01205	0.01341
		99%(short)	18 (0)	16.6	1.08	0.03370	0.01201	0.01341
FF	t-distrib	95%(long)	44	83	0.53	1.08262	0.08317	0.09529
		95%(short)	38	83	0.46	0.63702	0.08003	0.09333
		99%(long)	6 (0)	16.6	0.36	0.19659	0.13608	0.14974
		99%(short)	2 (0)	16.6	0.12	0.02356	0.13326	0.14766
	Empirical	95%(long)	84	83	1.01	2.22085	0.06522	0.07636
		95%(short)	80	83	0.96	1.59067	0.06160	0.07339
		99%(long)	16 (2)	16.6	0.96	0.50010	0.10006	0.11105
		99%(short)	18 (2)	16.6	1.08	0.23228	0.09454	0.10575
JY	t-distrib	95%(long)	38	83	0.46	40.5585	1.94391	2.18764
		95%(short)	23	83	0.28	14.3285	1.88372	2.18190
		99%(long)	9 (0)	16.6	0.54	10.2198	3.33207	3.60009
		99%(short)	3 (0)	16.6	0.18	1.56617	3.29305	3.59641
	Empirical	95%(long)	73	83	0.88	65.4063	1.51648	1.76698
		95%(short)	105	83	1.27	47.6068	1.28946	1.60556
		99%(long)	18 (1)	16.6	1.08	19.9317	2.75617	3.00797
		99%(short)	22 (4)	16.6	1.33	12.5997	1.95230	2.23977

Observed exceedences in the last 250 days are in brackets. MAD is mean absolute deviation.

Table 3.12. Likelihood Ratio Test Statistics of the Selected t-distribution Forecasts

Forecast Dist		Selected t-distribution			Empirical		
Test		Unconditional	Independence	Conditional	Unconditional	Independence	Conditional
Lower Tail	BP						
	95%	1.707	4.395	4.87*	0.049	2.550	2.599
	99%	0.070	N/A	N/A	0.145	N/A	N/A
	GM						
	95%	0.140	0.211	0.351	0.022	0.272	0.294
	99%	0.699	1.982	2.681	0.010	3.044*	3.053
	CD						
	95%	0.981	0.566	1.547	0.531	1.740	2.271
	99%	0.004	N/A	N/A	0.004	N/A	N/A
	FF						
	95%	0.022	0.446	0.468	0.087	0.202	0.289
	99%	0.287	0.603	0.889	0.050	0.752	0.803
	JY						
	95%	1.317	0.685	2.002	0.050	0.150	0.200
	99%	3.829*	1.066	4.895	0.145	N/A	N/A
Upper Tail	BP						
	95%	0.696	0.108	0.804	0.087	0.202	0.289
	99%	0.619	N/A	N/A	0.010	N/A	N/A
	GM						
	95%	0.089	0.468	0.558	0.342	0.101	0.444
	99%	0.004	N/A	N/A	0.010	N/A	N/A
	CD						
	95%	0.022	0.0	0.022	1.026	0.423	1.449
	99%	0.070	1.030	1.100	0.004	2.835	2.839
	FF						
	95%	0.203	0.046	0.249	0.006	0.143	0.148
	99%	0.472	N/A	N/A	0.010	N/A	N/A
	JY						
	95%	1.707	0.776	2.483	2.919*	0.001	2.919
	99%	0.370	N/A	N/A	0.472	N/A	N/A

At 5%, $\chi^2(1) = 3.84$ $\chi^2(2) = 5.99$. At 10% $\chi^2(1) = 2.71$ $\chi^2(2) = 4.61$ The highlighted likelihood ratio statistics are significant at 5% and those with * are significant at 10%.

Table 3.13. Likelihood Ratio Test Statistics of the Estimated t-distribution Forecasts

Forecast Dist		Estimated t-distribution			Empirical		
Test		Unconditional	Independence	Conditional	Unconditional	Independence	Conditional
Lower Tail	BP						
	95%	8.930	0.791	9.721	0.135	4.070	4.206
	99%	4.900	N/A	N/A	0.004	0.837	0.841
	GM						
	95%	8.407	0.723	9.130	0.087	1.048	1.135
	99%	6.042	N/A	N/A	0.004	2.835	2.839
	CD						
	95%	13.16	N/A	N/A	0.888	1.447	2.335
	99%	9.061	N/A	N/A	0.287	N/A	N/A
	FF						
	95%	10.04	0.939	10.97	0.005	1.297	1.303
	99%	3.933	N/A	N/A	0.010	0.929	0.939
	JY						
	95%	13.85	1.485	15.34	0.573	0.937	1.510
	99%	1.831	N/A	N/A	0.050	N/A	N/A
Upper Tail	BP						
	95%	14.57	0.016	14.58	0.135	0.171	0.306
	99%	9.061	N/A	N/A	0.189	N/A	N/A
	GM						
	95%	11.85	0.0	11.85	0.049	0.257	0.306
	99%	7.404	N/A	N/A	0.004	N/A	N/A
	CD						
	95%	13.16	0.442	13.60	2.053	0.187	2.240
	99%	N/A	N/A	N/A	0.287	0.603	0.889
	FF						
	95%	13.85	0.008	13.86	0.050	0.003	0.053
	99%	9.061	N/A	N/A	0.050	N/A	N/A
	JY						
	95%	27.46	N/A	N/A	2.468	.217	2.686
	99%	7.404	N/A	N/A	0.699	N/A	N/A

At 5%, $\chi^2(1) = 3.84$ $\chi^2(2) = 5.99$. At 10% $\chi^2(1) = 2.71$ $\chi^2(2) = 4.61$. The highlighted likelihood ratio statistics are significant at 5% and those with * are significant at 10%.

Table 3.14. Standard Deviation of the Forecast Variances under the Different Likelihoods.

Likelihood	BP	DM	CD	FF	JY
Normal	0.10296	0.04147	0.01900	0.04895	0.14800
Selected df	0.09595	0.03669	0.01815	0.03999	0.12417
Estimated df	0.11148	0.04215	0.01728	0.07142	0.13470

Table 3.15. Average Confidence Intervals of Forecast Empirical Distributions

Likelihood	Percentile	BP	DM	CD	FF	JY
Normal	5th	-1.60	-1.66	-1.58	-1.61	-1.61
	95th	1.63	1.64	1.61	1.64	1.44
	1st	-2.64	-2.75	-2.64	-2.62	-3.09
	99th	2.50	2.57	2.65	2.55	2.27
Selected t-dist.	5th	-1.65	-1.71	-1.64	-1.69	-1.75
	95th	1.69	1.69	1.67	1.72	1.56
	1st	-2.72	-2.85	-2.70	-2.73	-3.34
	99th	2.58	2.67	2.74	2.68	2.43
Estimated t-dist	5th	-1.57	-1.60	-1.58	-1.55	-1.61
	95th	1.61	1.57	1.59	1.56	1.41
	1st	-2.59	-2.65	-2.58	-2.49	-3.08
	99th	2.46	2.48	2.62	2.45	2.22

Fig. 3.1 Empirical Distributions of the Standardized Errors of the GARCH/t-distributions

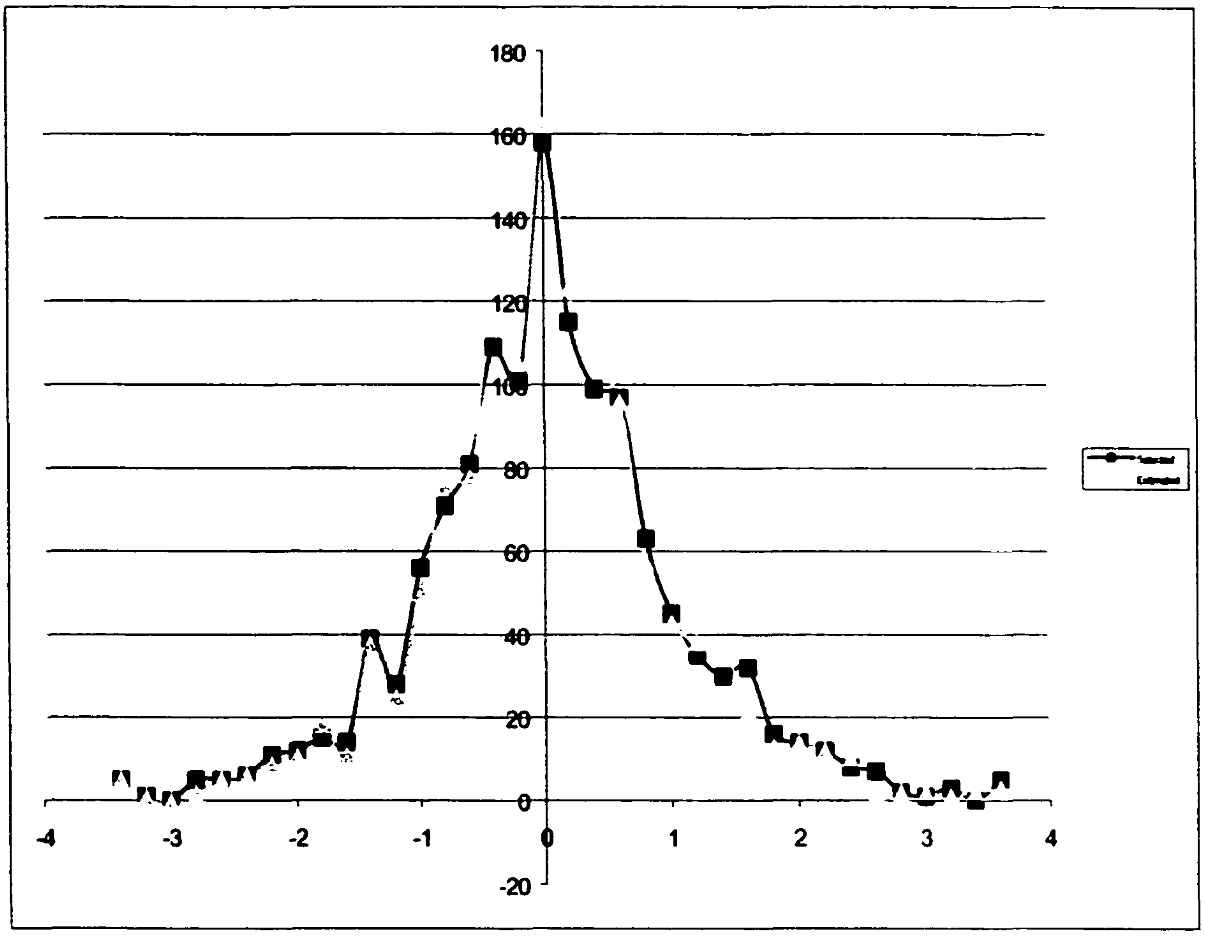

Reference

Akgiray, V. and Booth, G. G. (1988) " Mixed Diffusion-Jump Process Modeling of Exchange Rate Movements", *The Review of Economics and Statistics*, 70, 631-637.

Amin, K. I. And Ng, V. K. (1997) "Inferring Future Volatility from the Information in Implied Volatility in Eurodollar Options: A New Approach", *Review of Financial Studies*, 10(2),

Alexander, C. O. and Leigh, C. T. (1997) "On the Covariance Matrices used in Value At Risk Models", *Journal of Derivatives*, Spring, 50 - 62.

Baillie, R. T. and Bollerslev, T. (1989) "The Message in Daily Exchange Rates: A Conditional-Variance Tale." *Journal of Business & Economic Statistics*, 7(3), 297-305.

Barone-Adesi, G., K. Giannopoulos, and L. Vosper (1999) "VaR without Correlations for Portfolios of Derivative Securities," *Journal of Futures Markets*, 199-212.

Bates, D. S. (1996) "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options," *Review of Financial Studies*, 9(1), 69-107.

Bollerslev, T. R. (1986) "Generalized Autoregressive Conditional Heteroskedasticity," *Journal of Econometrics*, 31, 307-327.

______ (1987) "A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return," *The Review of Economics and Statistics*, 69(3), 542-547.

______, R. Y. Chou, and K. F. Kroner (1992) "ARCH Modeling in Finance," *Journal of Econometrics*, 52, 5-59.

______ and J.M. Wooldridge (1992) "Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time Varying Covariances." *Economics Reviews.*

Brooks, C. (1996) "Linear and Non-linear Forecastability of High-frequency Exchange Rate" *Journal of Forecasting*, 16(2), 125-146.

Calderon-Rossell, R. Jorge, and M. Ben-Horim (1982) "The Behaviour of Foreign Exchange Rates," *Journal of International Business Studies*, 13, 99-111.

Canina, L. and Figlewski, S. (1993) "The Information Content of Implied Volatility," *Review of Financial Studies*, 6(3), 659-682.

Christensen, B. J., and Prabhala, N. R. (1998) "The Relation between Implied and Realized Volatility." Journal of Financial Economics, 50, 125-150.

Christoffersen, P. F. (1998) "Evaluating Interval Forecasts," *International Economic Review*, 39(4), 841-862.

Dave, R. and S. Gerhard (1997) "On the Accuracy of VaR Estimates Based on the Variance-Covariance Approach," Working Paper, Olsen & Associates Research Institute, Zurich.

Day, T. E and Lewis, C. M. (1993) "Forecasting Future Market Volatility", *The Journal of Derviatives*, Winter 1993.

Duffie, D. and Pan, J. (1997) "An Overview of the Value at Risk", *The Journal of Derviatives*, Spring 1997.

El-Jahel, E., Perraudin, W., Sellin, P. "Value at Risk for Derivatives." *Journal of Derivatives*, Spring 1999.

Engle, R. F. (1982) "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation," *Econometrica*, 50, 987-1007.

______ and T. Bollerslev (1986) "Modeling the Persistence of Conditional Variance," Econometric Reviews, 5, 1-47.

______ and G. Gonzalez-Rivera (1991) "Semiparametric ARCH Models," Journal of Business and Economic Statistics, 9, 345-359.

Garman, M. B. and S. W. Kohlhagan (1983) "Foreign Currency Option Values," *Journal of International Money and Finance*, 2, 231-237.

Glasserman, P (2000) "The Quest for Precision through Value-at-Risk" Financial Times, Mastering RISK 2.

Heish, D. A. (1989) "Modeling Heteroskedasticity in Daily Foreign-Exchange Rates." *Journal of Business & Economic Statistics*, 7(3), 307-317.

Jorion, P. (1995) "Predicting Volatility in the Foreign Exchange Market", *The Journal of Finance*, 50(2), 507-528.

______ (1988) "On Jump Processes in Foreign Exchange and Stock Market," manuscript, Columbia University, Graduate School of Business.

Kendall, M. G., and A. Stuart (1969) *The Advance Theory of Statistics*, Vol. 2, 2nd Edition, London: Charles Griffin.

Longin, F. (2000) "From Value at Risk to Stress Testing: the Extreme Value Approach" *Journal of Banking and Finance*, 24, 1097-1130.

Lopez, J. A. (1996) "Regulatory Evaluation of Value-at-Risk Models," mimeo, Federal Reserve Bank of New York.

McFarland, J. W., R. Pettit, and S. K. Sung (1982) "The Distribution of Foreign Exchange Rate Changes: Trading Day Effect and Risk Measurement," *Journal of Finance*, 37, 693-715.

McLeod, A. J., and W. K. Li (1983) " Diagnostic Checking ARMA Time Series Models Using Squared-Residual Autocorrelations," Journal of Time Series Analysis, 4, 269-273.

McNees, S. (1995) "Forecast Uncertainty: Can It Be Measured?" mimeo, Federal Reserve Bank of Boston.

Melino, A., and S. M. Turnbull (1990) "Pricing Foreign Currency Options with Stochastic Volatility," *Journal of Econometrics*, 45, 239-265.

Rogalski, R. J. and J. D. Vinso (1978) "Empirical Properties of Foreign Exchange Rates," *Journal of International Business Studies*, 9, 69-79.

Westerfield, J. M. (1977) "Empirical Properties of Foreign Exchange Rates Under Fixed and Floating Exchange Rate Regimes," *Journal of International Economics*, 7, 181-200.

Xu, X. and S. J. Taylor (1995) "Conditional Volatility and the Informational Efficiency of the PHLX Currency Options Market," *Journal of Banking & Finance*, 19, 803-821.

CHAPTER 4

CONCLUSION

The thesis consists of two papers analyzing various aspects of return variability. In the chapter 2, we study the effects of extending trading time on trading activity, daily return variability, transitory variability in opening versus closing prices, trading versus non-trading time return variance, and the intraday return variability and volume patterns. We find that increasing trading time can generate trade if firms are opened during the additional period and information arrival about stocks is high. This suggests that continuous trading may not generate trading activity since most firms are closed during the overnight period and information flow about them is low. The study also finds that longer trading hours would not lead to an increase in daily return variability. This suggests that extending trading hours may not affect the risk premiums on stocks. The study of the effect of the early opening on intraday return variances sheds light on some of the hypotheses proposed in the microstructure literature. First, we find that trading time return variance did not increase relative to non-trading time return variance after adjusting for the effect of shift in time on these variances. This finding is consistent with the private information explanation of why trading time return variability is higher than non-trading time return variability. Second, we find that transitory variability in opening prices reduces relative to the transitory variability in closing prices after the early opening. This result supports the price formation hypothesis in explaining the observed higher transitory variance in opening prices.

In chapter 3, we study the roles of modeling unconditional kurtosis and the division of the kurtosis between the assumed distribution and time-varying return variances in accurately predicting the value at risk. We find that modeling kurtosis as measured by the number of standard deviations associated with a given level of confidence produce more accurate VaR forecasts. On the other hand, modeling kurtosis as traditionally measured may not improve VaR forecasts. This finding suggests that, for VaR of direct exposures, a more appropriate measure of kurtosis is the number of standard deviations associated with the particular confidence interval. The study also shows that, apart from modeling the unconditional kurtosis, its distribution between the imposed risk structure and time-varying variances is important in accurately forecasting some risk factors. For example, the study shows that appropriate fat-tail distributions rather than more general volatility models can properly model the BP exchange rate. This finding is important because it suggests that progress can be made in accurately forecasting the BP exchange rate risk by focusing research on identifying appropriate fat-tail distributions.