



A University of Sussex PhD thesis

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details

Returns, Volatility and Durations of High Frequency Financial Data

by

Xiufeng Yan

MSc, University of Sussex, 2013

Bachelor of Science, Southwestern University of Finance and Economics, 2012

This thesis is submitted in partial fulfillment of the
requirements for the degree of
Philosophy Doctor (Ph.D) in Mathematics

in the

School of Mathematical and Physical Sciences
University of Sussex

© Xiufeng Yan, 2020

Abstract

This dissertation consists of three chapters investigating the modelling of financial tick-by-tick data. Financial research using high-frequency data have been very active during the last two decades. The financial mathematical modelling of the high frequency price dynamics is based on the proper interpretation of the characteristics of the tick-by-tick data.

The first chapter provides an empirical investigation of the tick-by-tick returns. First, I provide a sampling method of returns different from the sampling from identical time interval. Second, I compare the returns from the two sampling methods using the Central Limit Theorem. The empirical results suggest that sampling returns from identical number of tick-by-tick transactions could recover the normality of intraday returns at lower costs due to its faster convergent rate.

The second chapter proposes a multiplicative component intraday volatility model. The intraday conditional volatility is expressed as the product of intraday periodic component, intraday stochastic volatility component and daily conditional volatility component. I extend the multiplicative component intraday volatility model of Engle (2012) and Andersen and Bollerslev (1998) by incorporating the durations between consecutive transactions. The model can be applied to both regularly and irregularly spaced returns. I also provide a nonparametric estimation technique of the intraday volatility periodicity. The empirical results suggest the model can successfully capture the interdependency of intraday returns.

The third chapter explores the duration dynamics modelling under the Autoregressive Conditional Durations (ACD) framework (Engle and Russell 1998). I test different distributions assumptions for the durations. The empirical results suggest unconditional durations approach the Gamma distributions. Moreover, compared with exponential distributions and Weibull distributions, the ACD model with Gamma distributed innovations provide the best fit of SPY durations.

Key Words: tick-by-tick data, Intraday volatility, Intraday seasonality, marked point process, UHF-GARCH models, intraday returns, Autoregressive Conditional Duration models, realized volatilities.

Acknowledgements

I am seriously in debt to so many people during the last five years in University of Sussex. It is fortunate for me that I can finally be able to say Thank you to all of you.

First, I would like to express my gratitude to my main supervisor, Qi Tang for guiding me through this process. There are some moments that have been particularly difficult for me. Nevertheless, Dr. Tang's constant patience, comprehension and support lead me through those dark times. I thank him for offering me his detailed comments and insights. When I struggled with doubt, his generous encouragement and trust helped me to move forward. I also thank him for giving me financial support for living cost and my participation to conferences. I might never be able to repay the extremely valuable help from him. Words cannot express my gratitude.

Second, I am extremely grateful to Professor Enrico Scalas. He patiently guided me to the world of statistical testing and provided sound advice that shaped my advance in the statistics. The third chapter regarding durations is directly based on our collaboration. I also thank him for inspiring my interests in ACD models.

I also thank all my professors in the PhD program who have helped me in various ways. A special thanks goes to Professor Sebastian Jaeger for his exceptional empathy and kindness. He granted me the extension so that I can finish this dissertation. I also thank him for his helpful suggestions. Besides, I am very thankful for Professor Anotida Madzvamuse. He has been my interviewer of the annual review for the last five years. He always could provide experienced advices and practical guidance regarding the development of my research.

I am grateful to my parents for always supporting and encouraging me. Their love and concerns constantly encourage me to face the difficult challenges bravely. In the end, I want to say thank you to my lovely neighbour Anqi who has given me a new sense of life. She is so considerate and selfless and never cease to encourage me. She made so many sacrifices so that I can complete the journey.

Table of content

Abstract.....	ii
Acknowledgements	iii
Table of content	iv
List of Tables.....	vi
List of Figures.....	viii
Chapter 1	1
1.1 Introduction.....	1
1.2 Sampling returns from time and sampling returns from transactions	3
1.2.1 The tick-by-tick data.....	3
1.2.2 Sampling Methodology	6
1.2.2.1 Literature Review	6
1.2.2.2 Time-aggregated returns and transaction-aggregated returns.....	8
1.3 Aggregational Normality	16
1.4 Tails of the return distributions.....	22
1.5 Realized volatility and standardized returns.....	28
1.6 Conclusions.....	35
Chapter 2	36
2.1 Introduction.....	36
2.2 Literature Review	38
2.3 Data	40
2.3.1 Descriptive statistics	41
2.3.2 Intraday return periodicity	43
2.4 Multiplicative component intraday volatility model.....	48
2.4.1 Interpretation of the intraday return volatility	48
2.4.2 Multiplicative component intraday volatility	51
2.5 Intraday volatility components and model estimation	55
2.5.1 ARCH effect of daily returns	55
2.5.2 Intraday volatility periodicity	64
2.5.3 Econometrics regarding estimation.....	66
2.6 Empirical results.....	68
2.6.1 Normalized intraday returns and filtered intraday returns.....	69
2.6.1.1 Normalized intraday returns	69
2.6.1.2 The filtered intraday returns.....	73
2.6.2 Results of GARCH estimation and volatility persistence	77
2.7 Conclusion.....	81
Chapter 3	83
3.1 Introduction.....	83
3.2 ACD framework	85
3.2.1 Theoretical background of ACD.....	85
3.2.1.1 Point Process.....	85

3.2.1.2 ACD models	87
3.2.1.2.1 Duration specification	89
3.2.1.2.2 Distributional assumption and maximum likelihood estimation	92
3.2.1.3 Durations and marks.....	97
3.2.1.3.1 General setup	97
3.2.1.3.2 Durations and Volatility.....	99
3.2.1.3.3 Durations and other explanatory variables.....	103
3.2.2 ACD in applications	107
3.2.2.1 Intraday Periodicity	108
3.2.2.2 Testing the ACD.....	111
3.2.2.3 ACD papers	115
3.2.2.3.1 Trade durations	115
3.2.2.3.2 Price durations.....	119
3.2.2.3.3 Volume durations and others	121
3.3 Data and some empirical results.....	122
3.3.1 Market microstructural issues and transaction-aggregated durations.....	122
3.3.2 The descriptive statistics of transaction-aggregated durations	125
3.3.3 Intraday seasonality of durations	128
3.3.4 Unconditional distributional properties of the durations	134
3.4 ACD modelling of durations.....	146
3.4.1 Data preparation: Intraday seasonality and re-initialization	146
3.4.2 ACD modelling of transaction-aggregated duration dynamics.....	150
3.5 Conclusions.....	156
Appendix A.	157
Appendix B.	159
Appendix C.	159
Appendix D.	159
References.....	162

List of Tables

Table 1.1 Descriptive statistics of tick-by-tick returns and tick-by-tick durations.....	4
Table 1.2 Descriptive statistics of time-aggregated returns	10
Table 1.3 Descriptive statistics of transaction-aggregated returns	10
Table 1.4 Autocorrelation of time-aggregated return series.....	13
Table 1.5 Autocorrelation of transaction-aggregated return series.	14
Table 1.6 Autocorrelation of time-aggregated squared return series	15
Table 1.7 Autocorrelation of transaction-aggregated squared return series	15
Table 1.8 EDF statistics of time-aggregated returns	17
Table 1.9 EDF statistics of transaction-aggregated returns.....	17
Table 1.10 Critical values of EDF statistics for the null hypothesis of normal distribution .	17
Table 1.11 Summary of within-day normality test on time-aggregated returns.....	20
Table 1.12 Summary of within-day normality test on transaction-aggregated returns.	20
Table 1.13 The modelling of tails of time-aggregated returns using generalized pareto distribution	26
Table 1.14 The modelling of tails of transaction-aggregated returns using generalized pareto distribution	26
Table 1.15 The summary of the realized volatilities of time-aggregated returns.....	31
Table 1.16 The summary of the realized volatilities of transaction-aggregated returns	31
Table 1.17 Descriptive statistics of standardized time-aggregated returns at different frequencies	32
Table 1.18 Descriptive statistics of standardized transaction-aggregated returns at different frequencies	32
Table 1.19 Summary of within-day normality test on standardized returns	33
Table 1.20 Summary of SPY daily returns.....	34
Table 1.21 Summary of SPY daily realized variance.....	34
Table 2.1 Descriptive statistics of time-aggregated returns	41
Table 2.2 Descriptive statistics of 5-minute return series, absolute 5-minute return series and squared 5-minute return series.	41
Table 2.3 Descriptive statistics of 10-minute return series, absolute 10-minute return series and squared 10-minute return series.....	42

Table 2.4 Autocorrelations of average returns and absolute average returns.....	46
Table 2.5 Descriptive statistics of daily return series, absolute daily return series and squared daily return series	55
Table 2.6 Parameter estimation of the MA (1)-GARCH (1,1) model.....	59
Table 2.7 Correlations between GARCH forecasts and other ex post daily variability measures.....	63
Table 2.8 The normalized time-aggregated returns.....	70
Table 2.9 The normalized transaction-aggregated returns	70
Table 2.10 The filtered time-aggregated returns	74
Table 2.11 ARCH test of filtered time-aggregated returns	75
Table 2.12 GARCH estimation of intraday time-aggregated returns	78
Table 2.13 GARCH estimation of filtered intraday time-aggregated returns	79
Table 3.1 Descriptive statistics of transaction-aggregated durations	125
Table 3.2 Descriptive statistics of tick-by-tick durations with different length	137
Table 3.3 The fitting of Exponential, Weibull, Gamma and Generalized Pareto distribution using durations less than 1s	139
Table 3.4 The fitting of Exponential, Weibull, Gamma and Generalized Pareto distribution using durations with length between 1s and 60s	140
Table 3.5 The fitting of Exponential, Weibull, Gamma and generalized Pareto distribution to transaction-aggregated durations.....	142
Table 3.6 The fitting of Exponential, Weibull, Gamma and generalized Pareto distribution to transaction-aggregated durations.....	143
Table 3.7 The fitting of Exponential, Weibull and Gamma distributions to within-day transaction-aggregated duration	144
Table 3.8 Descriptive statistics of the raw transaction-aggregated durations	149
Table 3.9 Descriptive statistics of the filtered transaction-aggregated durations.....	149
Table 3.10 ACD estimation of T-13 transaction-aggregated durations.....	151
Table 3.11 ACD estimation of T-67 transaction-aggregated durations.....	152
Table 3.12 ACD estimation of T-134 transaction-aggregated durations.....	153
Table 3.13 ACD estimation of T-400 transaction-aggregated durations.....	153

List of Figures

Figure 1.1 SPY tick-by-tick returns	4
Figure 1.2 SPY tick-by-tick waiting-time	5
Figure 1.3 Positive tail of SPY 5-minute returns	28
Figure 1.4 Negative tail of SPY 5-minute returns	28
Figure 2.1 The average SPY 5-minute return over the trading day	43
Figure 2.2 The average SPY 10-minute return over a trading day	44
Figure 2.3 The average SPY 10-minute return over a trading day	45
Figure 2.4 The average SPY 10-minute return over a trading day	45
Figure 2.5 Five days SPY absolute 5-minute return correlogram	47
Figure 2.6 Five days SPY absolute 10-minute return correlogram	47
Figure 2.7 SPY daily returns from January 3, 2005 to December 31, 2014	56
Figure 2.8 SPY absolute daily return correlogram	57
Figure 2.9 SPY squared daily return correlogram	57
Figure 2.10 The conditional variance and standardized residuals of MA (1)-GARCH (1,1) fitting	60
Figure 2.11 The correlogram of squared standardized residuals of MA (1)-GARCH (1,1) fitting	60
Figure 2.12 The Quantile-Quantile Plot of the standardized residuals	61
Figure 2.13 The absolute daily returns and GARCH (1,1) volatility forecasts	62
Figure 2.14 The cumulative absolute returns, realized volatility and GARCH (1,1) forecasts	63
Figure 2.15 One-day-ahead MA (1) -GARCH (1,1) estimation	69
Figure 2.16 The correlogram of normalized 5-minute absolute returns	72
Figure 2.17 The correlogram of normalized 10-minute absolute returns	73
Figure 2.18 The estimated intraday periodicity for each 5-minute time interval	74
Figure 2.19 The correlogram of filtered 5-minute absolute returns	76
Figure 2.20 The correlogram of filtered 10-minute absolute returns	76
Figure 2.21 Residuals correlogram of the filtered 5-minute returns	80
Figure 2.22 Residuals correlogram of the filtered 10-minute returns	80
Figure 2.23 The Q-Q plot of filtered 5-minute residuals	81

Figure 3.1 Correlogram of tick-by-tick durations	127
Figure 3.2 Correlogram of T-13 durations	127
Figure 3.3 Correlogram of T-67 durations	128
Figure 3.4 The tick-by-tick duration of SPY	128
Figure 3.5 The average tick-by-tick duration of SPY over the trading day	129
Figure 3.6 The correlogram of SPY average tick-by-tick durations	130
Figure 3.7 The correlogram of T-800 transaction-aggregated durations.....	131
Figure 3.8 The correlogram of T-400 transaction-aggregated durations.....	132
Figure 3.9 Cubic spline function with 15-minute nodes of SPY tick-by-tick duration.....	133
Figure 3.10 The kernel density function of tick-by-tick durations.....	134
Figure 3.11 The kernel density function of T-134 durations.....	135
Figure 3.12 The kernel density function of T-400 durations.....	135
Figure 3.13 The kernel density function of T-800 durations.....	135
Figure 3.14 The ECDF of tick-by-tick durations less than 1s	138
Figure 3.15 The ECDF of tick-by-tick durations with length between 1s and 60s	139
Figure 3.16 The histogram of tick-by-tick durations longer than 60s	141
Figure 3.17 The empirical cumulative distribution functions of T-24000 transaction- aggregated durations.....	143
Figure 3.18 The empirical cumulative distribution functions of T-40000 transaction- aggregated durations.....	144
Figure 3.19 Cubic spline function with 15-minute nodes of SPY tick-by-tick durations ...	148
Figure 3.20 The correlogram of T-400 GACD (1,1) residuals.....	155
Figure 3.21 The histogram of T-400 GACD (1,1) residuals	155

Chapter 1

High frequency returns: The comparison of two sampling methods.

1.1 Introduction

The availability of large data computing techniques has reached a stage that allows for the investigation of a wide range of issues in the empirical high frequency financial research.

According to the work of Goodhart and O'Hara (1997), empirical financial research using high frequency return data can be classified into three main categories.

The first group studies the effects of market structure and its relationship with the interpretation and availability of high-frequency data. For data recorded at the finest level of frequency, it is hoped that the characteristics of trading mechanism and operational details and their effects on the market behavior can be revealed. For instance, discussions with respect to the minute operational details of the NYSE can be found in Hasbrouck et al. (1993) and O'Hara (1995); discussions on the price limits and mandatory trading halts in TSE has been presented in Lehmann and Modest (1994).

The second group focuses on the statistical description and dynamic modelling of different random variables associated with the high frequency data set such as prices, volumes and waiting-time. An important feature of tick-by-tick data is that the distances between observations are not equally spaced in time. Thus, the waiting-time that is the time differences between two consecutive transactions can also be viewed as a random variable. This feature naturally suggests the use of marked point process in the modelling of waiting-time series since point process could jointly model the arrivals of event and other information (marks) associated with events. Engle and Russell (1997) advocated the autoregressive duration model for the modelling of waiting-time between transactions.

Scalas et al. (2002) discussed the use of continuous time random walk model for the description of tick-by-tick dynamics. There are many papers with respect to the discussions concerning high frequency volatility: the use of integrated volatility that is defined as the sum of high frequency squared returns are suggested by Andersen et al. (1999), Bollerslev and Wright (2000) and Barndorff-Nielsen and Shephard (2002); modifications of the classic GARCH volatility model for intraday volatility dynamics can be found in Engle and Sokalska (2011), Russell (1999) and Bollerslev (1996).

The third group investigates the information content of high frequency data and explores the possibility of using high frequency data to analyze the limitation of market efficiency. Their focus is on the analysis of the information flows and its relationship with price movements. For instance, Barclay and Warner (1993) argued that the medium sized trades have the highest possibility of being informed trades. Easley et al. (1996) also used intra-day data to test the size effects. Jones et al. (1994) and Lyons (1994) argued that the transaction intensity carries more information than volume of transactions. Research on such effects is commonly known as the market micro-structure studies.

A main issue regarding financial research using high frequency data is the sampling method of the returns. The statistical and economic conclusion drawn from high frequency data is directly based on the sampling method of returns. Given the nature of tick-by-tick records that the occurrence of transactions is irregularly spaced in time, it is of both practical and theoretical importance to consider sampling method other than calculating returns from fixed time intervals. Moreover, the unconditional return distribution is of both theoretical and practical importance in many pricing and risk management related areas.

In this chapter, I focus on the unconditional distributional properties of the tick-by-tick returns of SPY. I first discuss the nature of time in sampling high frequency returns and explain my sampling methodology. I then compare my sampling method with sampling from fixed time interval and discuss some stylized facts in empirical research using high frequency data. I provide a return sampling method that could reflect the variability of asset price in the busy trading hours and show this sampling method could recover the normality

to the return series. I hope the data analysis of an equity with extremely high liquidity could contribute to the understanding of ultra-high-frequency financial data and form the bases for the next two chapters. Chapter 1 is organized as follows:

Section 1.2 is the introduction of data and our sampling methodology. It also consists of analysis of the statistical characteristics of the return series. Section 1.3 analyzes the aggregation normality. Section 1.4 exhibits the discussion regarding the tail properties of the high frequency returns. Section 1.5 discusses the realized volatility and the standardized returns. Section 1.6 is the conclusion.

1.2 Sampling returns from time and sampling returns from transactions

In this section, I briefly introduce my data and explain my sampling methodology.

Section 1.2.1 is the introduction of the tick-by-tick data. Section 1.2.2 presents the sampling methodology.

1.2.1 The tick-by-tick data

The data set is the tick-by-tick transaction records of SPDR S&P 500 ETF Trust (SPY) during 01/02/2014-12/31/2014. The data is collected from the company TICK DATA. SPY is chosen since it reflects the overall movement of the market and it is one of the most frequently traded equities. The term of “tick-by-tick” means the most detailed display of the trading information. Specifically, every trade is contained in the data. It contains every transaction for 252 trading days. The out-of-hour transactions outside the time period between 9:30 am and 4:00 pm on each day are removed. There are 79156264 tick-by-tick transaction records.

The data set consists of prices $p_{i,d}$, volumes $v_{i,d}$ and execution time $t_{i,d}$. Where

- d is the day index, varying from 1 to the total number of trading days, 252.
- i is the trade index, varying from 1 to the total number of trades occurring in day d .

I define the waiting-time $w_{i,d}$ as $w_{i,d} = t_{i+1,d} - t_{i,d}$. The tick-by-tick returns $r_{i,d}$ is defined as $r_{i,d} = \ln \left(\frac{p_{i+1,d}}{p_{i,d}} \right)$.

	Sample Size	Mean	Variance	Kurtosis	Skewness	Maximum	Minimum
Tick-by-Tick returns	79156012	3.2285E-10	7.30409E-09	60236.25405	-0.600116	0.0378838	-0.0379332
Tick-by-Tick Waiting-times	79156012	0.074495162	0.211414975	3407788.065	1085.8828	1808.476	0

Table 1.1 Descriptive statistics of tick-by-tick returns and tick-by-tick durations.

Table 1.1 gives a brief statistical description of the tick-by-tick return and tick-by-tick waiting-time series. With respect to the tick-by-tick return series, the mean value of the returns is nearly indistinguishable from zero. In fact, there are many zero returns that result in the extremely high kurtosis. The extremely high kurtosis and negative skewness suggest that the distribution of the returns is obviously not normal. The absolute values of the maximum and minimum of tick-by-tick returns are very close to each other.

With respect to the tick-by-tick waiting-time series, the waiting-time is counted in seconds and the mean value of the duration between two consecutive tick-by-tick transactions is around 0.07s. The mean indicates that the SPY is very frequently traded. The variance of the series is around 0.21 seconds squared. The maximum of the waiting-time is around 1808s, meanwhile the minimum of the waiting-time series is 0s. The minimum of 0s for tick-by-tick waiting-time series represents that there are multiple transactions executed at the same time. Detailed analysis regarding the waiting-time series of SPY is presented in Chapter 3.

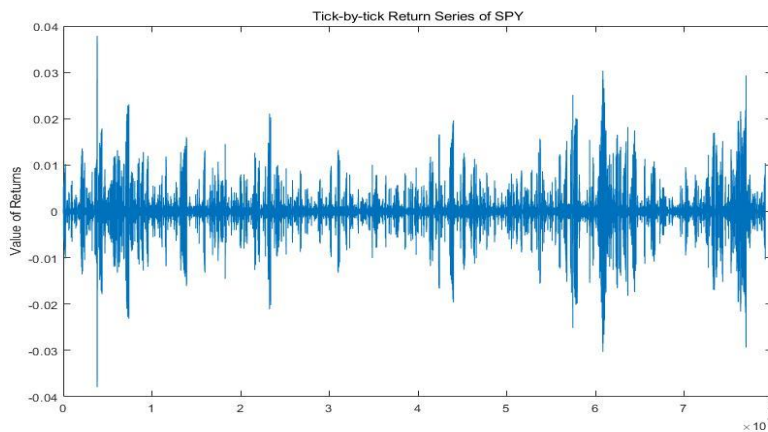


Figure 1.1 SPY tick-by-tick returns

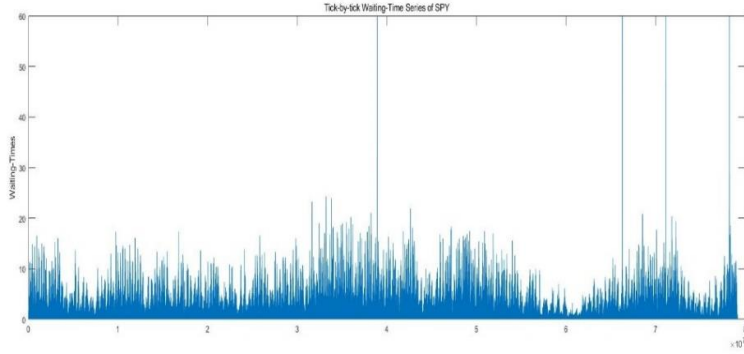


Figure 1.2 SPY tick-by-tick waiting-time

I conjecture from Table 1.1 that the operational details and transaction recording mechanism play a decisive role in the formation of the data pattern for a symbol like SPY with high liquidity. Unfortunately, these knowledges cannot be easily obtained and therefore their influences on the data pattern are hard to capture and to exclude. A direct use of the tick-by-tick transaction records might suffer from the distortion induced by such market microstructural details and further rises many difficulties in the construction of price dynamic model. For instance, since the average value of the tick-by-tick waiting-time series is around 0.07s, it is very difficult to capture the influence of intraday volatility periodicity on the price dynamics during such a short time interval since the estimation requires the full historical trading information in that short time interval, which generally causes large computation difficulties and is elaborated in Chapter 2. Also, from the economic perspective, it is highly unlikely that traders can respond any meaningful information in such a short time interval. Although computer trading programs could respond in time for market price movement information in time, it is hard to capture the trading mechanism behind such programs.

In order to draw meaningful results from both economic and statistical perspective, I need to apply some standard for the aggregation of the tick-by-tick data. The question is that which standard can be used to aggregate the tick-by-tick records and what can be learned from such sampling method. I discuss the issue of sampling methodology in the next section.

1.2.2 Sampling Methodology

In the construction of sampling method, I apply the “data-based” approach. One important principle is: “let the data speak for themselves”. It is argued by Engle (2006), “in addition to a model’s capability of producing statistical property that is observed in high frequency time series, the ultimate way to evaluate a model is its usefulness in many relevant areas such as option pricing and portfolio construction”. To circumvent the difficulties caused by the market microstructure effect and to avoid the unnecessary discard of data, I aggregate the tick-by-tick data by random variables associated with the transaction records. The natural candidates are volumes, transaction numbers and time since these are available random variables in the data. I further discuss those variables and their relations with price dynamics in the following literature review.

1.2.2.1 Literature Review

The dependence of the price process on volumes has been recognized for a long time. Clark (1973) argues that the variation of trading volumes may account for the deviation of returns from normal distribution. Karpoff (1987) further suggests that in equity markets, volume rises more when prices are rising than when prices are falling. The dependence of the prices process on volume implies that volatility will also be affected by volume. Since then, the relationship between trading volume and stock prices has been widely investigated theoretically and empirically (Gallant et al. 1992, Campbell et al. 1993, Lastrapes and Lamoureux 1990). Almost all of these empirical researches suggest that the volume is positively correlated with volatilities.

Another interesting empirical finding is, as summarized by Goodhart and O’Hara in 1997, that the volume is loosely, negatively correlated with the waiting-time. They further argued that where the price process goes will differ depending upon whether volume is high or low. Easley et al. (1994) analyzed the effect of the trade size on the price movement. They find that it is the occurrence of the transactions, rather than volume that moves the market. Jones et al. (1994) also suggested that volume has significantly less explanatory power than the number of transactions per second over the price movements. Such results imply that

it is the intensity of the occurrence of transactions, rather than the volume of those transactions, that moves the market since the number of transactions during fixed time interval is a measure of trading intensity.

The intuitive assumption behind the use of market data sampled from fixed time interval in both empirical and theoretical financial research is that whatever drives the price process, it should be relatively stable during a short time interval. Developments on the research of market microstructure that focus on the decision-rules followed by price-setting agents clearly challenge this assumption in the context of high frequency financing (Goodhart and O'Hara 1997).

A fundamental property of tick-by-tick data is that observations are not equally sampled in the time line. The “intra-day seasonal” such as the U shape of volumes and volatilities of equities on NYSE market (Lockwood and Linn 1990, Ghysels et al. 1993) makes the choice of optimal time interval very difficult. Specifically, if a short interval is chosen, there will be many intervals that contain no new information. In contrast, if we select a long interval, then information contained in the interval might be lost. Moreover, as argued by Diamond and Verrecchia (1987), short-sell constraint could impart information to no-trading intervals. Consequently, the waiting-time is not independent of the prices process. The result of Jone et al. (1994) further supported this implication by showing that the number of transactions is closely related to change of prices since transactions are more likely to happen when there is new information. Ane and Geman (2000) argue that the numbers of trades have more explanatory power over the price dynamics than the trading volumes. As a result, sampling by equal number of transactions should be more appropriate than sampling by equal amounts of trading volumes since it captures the volatility more accurately.

There are two main advantages of this sampling method. First, it retains the information of waiting-time series. Second, since data is sampled by equal number of transactions, it reflects the variability of prices during the most frequently traded hours more accurately. Based on previous literatures, I sample from the tick-by-tick transaction data by both fixed

time intervals and fixed number of transactions. I compare the two sampling methods by analyzing the statistical characteristics of the aggregated returns series at different frequencies and discuss some stylized facts in the next section.

1.2.2.2 Time-aggregated returns and transaction-aggregated returns

In this section I present the descriptive statistics of the return series from two different sampling methods at different frequencies. Recall that our data set consists of prices $p_{i,d}$, volumes $v_{i,d}$ and execution time $t_{i,d}$, where d is the day index, varying from 1 to the total number of trading days 252 and i is the trade index, varying from 1 to the total number of trades occurring in day d .

With respect to the returns sampled from fixed number of transactions, suppose that I use every consecutive T tick transaction records to calculate my non-overlap return series, then $(T)r_{i,d}$, the i th return on day d of the aggregated return series, is defined as

$$(T)r_{i,d} = \sum_{j=(T-1)(i-1)+1}^{(T-1)i} \ln \left(\frac{p_{j+1,d}}{p_{j,d}} \right) = \ln \left(\frac{p_{(T-1)i+1,d}}{p_{(T-1)(i-1)+1,d}} \right),$$

where d is the day index, varying from 1 to the total number of trading days 252, i is the trade index, varying from 1 to the nearest integer less than or equal to the total number of trades occurring on day d divided by T . I use the terminology “*transaction-aggregated returns*” to describe the returns calculated from fixed number of transactions. Specifically, by “*T-1000 return series*”, I mean the transaction-aggregated return series with $T = 1000$.

Regarding the returns sampled from fixed time intervals, suppose that I calculate the returns from every Δt minute, then for the time interval $[\tau, \tau + \Delta t]$ on day d , I define $s_{\tau,d}$, the price of SPY at τ on day d , by the average price of the two transactions whose occurrence are closest to τ . Specifically, $s_{\tau,d} = \frac{p_{m,d} + p_{n,d}}{2}$, where

$$m = \max \{i | t_{i,d} \leq \tau\} \text{ and } n = \min \{i | t_{i,d} \geq \tau\}.$$

$s_{\tau+\Delta t,d}$ are defined similarly. Then I use $\ln \left(\frac{s_{\tau+\Delta t,d}}{s_{\tau,d}} \right)$ to define the return during the time interval $[\tau, \tau + \Delta t]$ on day d . I use the terminology “*time-aggregated returns*” to describe

the returns calculated from fixed time intervals. Specifically, by “*1-minute return series*”, I mean the time-aggregated return series with $\Delta t = 1$ minute.

In order to fully exploit our tick-by-tick transaction data and to compare the return series measured at different frequencies, I use $\Delta t = 30s, 1min, 1.5min, 3min, 5min, 10min, 15min, 30min, 78min$ for the time-aggregated return series. The reason for not using tick-by-tick returns is presented in section 1.2.1.

I choose the values of T for the transaction-aggregated return series according to the sample sizes of time-aggregated returns to facilitate the comparison between the two sampling methods. For instance, since the 30-second return series has total 195834 observations, I choose $T=400$ accordingly to yield the $T=400$ return series with 198013 observations. In order to rule out the overnight effect, both time-aggregated and transaction aggregated returns are calculated within each trading day. Thus, it is difficult to choose the value of T to generate a transaction-aggregated return series with identical size of corresponding time-aggregated return series since transactions within a trading day varies.

The average and minimum of SPY tick-by-tick transactions occurred in one trading day are 314111 and 121489 respectively. The high liquidity of SPY allows me to use relatively large value of T to calculate the intraday transaction-aggregated returns. I use $T = 400, 800, 1200, 2400, 4000, 8000, 12000, 24000$ and 62000 for the transaction-aggregated returns.

	Sample Size	Mean	Variance	Kurtosis	Skewness	Maximum	Minimum
30Sec	196560	1.2961E-07	5.30347E-08	264.8727029	-0.173806698	0.01225414	-0.012983101
1Min	98280	2.5923E-07	9.67974E-08	95.91646403	-0.103161028	0.012458815	-0.012906394
1.5Min	65520	3.8885E-07	1.39691E-07	80.86530506	-0.005621653	0.012036628	-0.013264313
3Min	32760	7.7770E-07	2.70802E-07	39.85096976	-0.238157281	0.011012149	-0.014459855
5Min	19656	1.2961E-06	4.06658E-07	8.436150587	0.075305595	0.006182118	-0.005359357
10Min	9828	2.5923E-06	8.03275E-07	7.918815602	-0.037720742	0.008269006	-0.006320998
15Min	6552	3.8885E-06	1.17727E-06	7.388130566	-0.079263515	0.007751481	-0.006653523
30Min	3276	7.7770E-06	2.34033E-06	6.660810707	-0.095190193	0.011082252	-0.009269532

78Min	1260	2.0220E-05	6.20987E-06	6.687141886	-0.195275515	0.014103076	-0.014359004
-------	------	------------	-------------	-------------	--------------	-------------	--------------

Table 1.2 Descriptive statistics of time-aggregated returns

	Sample Size	Mean	Variance	Kurtosis	Skewness	Maximum	Minimum
T=400	198261	1.18123E-07	5.36498E-08	2878.43145	0.235143736	0.02944752	-0.029302461
T=800	98942	2.24205E-07	9.13855E-08	401.405859	0.053587616	0.01968218	-0.019585466
T=1200	65893	3.36582E-07	1.50566E-07	1046.22599	0.090477395	0.02954424	-0.029495885
T=2400	32876	1.6904E-06	2.66206E-07	345.1177319	5.896603684	0.02988284	-0.002934328
T=4000	19663	1.47527E-06	4.05294E-07	3.62073174	0.015212314	0.00385252	-0.003987089
T=8000	9763	2.7081E-06	8.23168E-07	3.51020343	0.00229587	0.00562988	-0.003523489
T=12000	6468	5.59325E-06	1.22472E-06	3.41251923	-0.039828733	0.00639776	-0.004232986
T=24000	3162	1.30382E-05	2.48051E-06	3.490401133	-0.041531592	0.00904297	-0.005727332
T=62000	1165	1.45991E-05	6.70171E-06	3.009298811	-0.054308977	0.00828105	-0.008328602

Table 1.2 Descriptive statistics of transaction-aggregated returns

With respect to the mean and variance, the sample means of transaction-aggregated return series and time-aggregated return series are nearly indistinguishable from zero. The means monotonically increase as the sample frequencies decreases. The variances of both transaction-aggregated return series and time-aggregated return series also increase monotonically as the sample frequencies decrease. The mean and variance of each time-aggregated returns are very close to the mean and variance of the corresponding transaction-aggregated returns.

In terms of the kurtosis, the fourth standardized moment, it decreases sharply with the decreasing return frequencies. For the transaction-aggregated returns, it decreases from 2878.43 of the *T-400* return series to 3.01 of the *T-62000* return series. For the time-aggregated returns, it drops from 264.87 of the 30-second return series to 6.69 of the 78-minute return series. An intriguing result is that transaction-aggregated return series have higher values of kurtosis when returns are measured at highest frequencies. E.g. The kurtosis of *T-400* return series and *T-800* return series are 2878.43 and 401.41 respectively. In contrast, the kurtosis of the corresponding 30-second and 1-minute return series are 264.87 and 95.92 respectively. Meanwhile time-aggregated returns exhibit higher kurtosis at relatively low frequencies. For instance, the kurtosis of 15-minute, 30-minute and 78-min return series are 7.39, 6.66 and 6.68 respectively. In comparison, the kurtosis of *T-12000*, *T-24000* and *T-62000* return series are 3.41, 3.49 and 3.01 respectively. Regarding this phenomenon, I present a plausible explanation.

Since the transaction-aggregated return series are sampled from fixed number of transactions, most of the observations in the series are calculated from tick-by-tick transactions during the busiest trading hours. As a result, the smaller the fixed number of transactions is used for the sampling, the higher the probability of a return observation being sampled from the busiest trading hours is. Thus, high frequency transaction-aggregated return series exhibit larger kurtosis since it mainly describes the highly volatile price movement during the trading hours with the highest intensity of transactions. If a larger number is used for aggregating the tick-by-tick returns, the Central Limit Theorem asserts that the transaction-aggregated returns should be approximately normal since each return in a transaction-aggregated return series is the sum of tick-by-tick returns. In contrast, returns in a time-aggregated return series are not the sum of an identical number of tick-by-tick returns. Although the Central Limit Theorem also applies to time-aggregated returns, the irregular number of tick-by-tick returns in each time interval might suggest that the time-aggregated returns should not be considered to follow some identical normal distribution.

With respect to the skewness that is the third standardized moment, transaction-aggregated return series and time-aggregated return series exhibit very different features. A distinct feature is that time-aggregated return series tend to have negative skewness while the transaction-aggregated returns tend to have positive skewness. For instance, almost all the time-aggregated returns exhibit negative skewness and the only exception is the 5-minute return series with a skewness of 0.075. In contrast, there are only three of the nine transaction-aggregated returns with negative skewness. Specifically, the skewness of $T-12000$, $T-24000$ and $T-62000$ are -0.040, -0.042 and -0.054 respectively.

Although the absolute value of skewness of transaction-aggregated returns tend to decrease with the decreasing return frequencies, it does not have such tendency for time-aggregated returns. For instance, the skewness of 30-second return series is -0.174 while the skewness of 78-minute return series is -0.195. Among time-aggregated returns, the 1.5-minute return has the smallest absolute value of skewness 0.006. Among transaction-aggregated returns, the $T-8000$ has the smallest absolute value of skewness 0.002.

The absolute values of maximum and minimum of both transaction-aggregated returns and time-aggregated returns fluctuate around 0.01. The maximum of transaction-aggregated returns generally decreases as the return frequencies decrease. Specifically, it drops from 0.029 for the *T-400* return series to 0.008 for the *T-62000* return series. It might be an implication that the abnormal returns can only occur in a very short time interval, which is in line with the theory of market efficiency since the theory asserts that most abnormal returns are associated with new information which will soon be common knowledge.

It is widely identified that price movements in liquid markets do not exhibit any significant autocorrelations when returns are measured at relatively low frequency. The absence of significant autocorrelations in returns are often considered as a support of the market efficiency. Discussions regarding the relationship between the absence of autocorrelations in returns and the market efficiency can be found in Fama (1971,1994) and Plerou et al. (1999).

However, the significant autocorrelation, usually negative, at very short lags is documented for most of the empirical financial research using high frequency data. Market microstructure research (Goodhart and O'Hara 1997 and Cambell et al. 1999) explain this phenomenon by the bid-ask bounce. Another explanation is the feedback trading that is the trading behavior conducted by investors who extrapolate from historical prices. See e.g. Shiller (1984), Culter et al. (1990), Sentana & Wadhwani (1992) and Merton (1980).

The volatility clustering is the phenomena that large price variations are more likely to be followed by large price variations. This phenomenon is commonly described by the positive autocorrelations of the squared returns. I now present the autocorrelations of the transaction-aggregated returns and time-aggregated returns. The significance level for the upper and lower bounds is 5%.

	30S	1Min	1.5Min	3Min	5Min	10Min	15Min	30Min	78Min
Lag 1	-0.11357	-0.08265	-0.06652	-0.06156	-0.01506	-0.02468	-0.01052	0.018744	0.078542
Lag 2	-0.00388	0.006217	-0.00964	-0.0166	-0.02727	-0.00717	-0.00204	0.058162	-0.03436

Lag 3	-0.00409	-0.01014	0.002306	-0.01246	0.011513	0.025502	0.043361	0.016442	-0.00045
Lag 4	0.006286	-0.00573	-0.00376	0.002623	-0.00224	-0.01868	0.035667	-0.00274	-0.02225
Lag 5	-0.00242	0.002863	-0.01425	0.013327	-0.00199	0.038865	-0.00593	0.003051	0.003938
Lag 6	-0.00057	-0.00204	-0.00986	-0.00241	0.007446	0.029326	0.010969	-0.04483	-0.00908
Lag 7	-0.00865	-0.00852	0.003313	-0.00739	0.005399	0.009408	0.022917	0.009573	-0.02074
Lag 8	-0.0007	-0.0142	-0.00011	0.01208	-0.00443	-0.00226	-0.01917	-0.00182	0.067125
Lag 9	-0.00036	-0.00551	0.009087	-0.00773	0.004174	-0.00252	-0.00845	-0.01438	0.031809
Lag 10	0.003057	0.005903	0.008632	-0.00018	0.023175	0.02309	0.039042	-0.02076	-0.02454
Lag 11	0.001128	-0.00305	-0.00708	0.006667	0.022194	0.008302	-0.02248	0.013483	-0.02821
Lag 12	0.000678	0.003842	-0.00192	0.01046	0.008997	-0.01527	-0.03088	0.01013	0.007252
Lag 13	-0.00305	0.00428	-0.00589	-0.01207	-0.00525	-0.01152	-0.01314	-0.01928	0.015616
Lag 14	-0.00631	-0.00137	0.004525	0.003802	0.013263	0.025103	0.014792	-0.0149	-0.01386
Lag 15	-0.0049	0.010166	-0.00576	0.003795	0.007467	0.009805	0.012154	0.002383	0.042766
Lag 16	-0.00601	0.002462	0.004921	-0.00159	-0.00964	0.004594	-0.01175	-0.01162	0.008348
Lag 17	-0.00726	-0.00316	0.004827	0.022055	-0.00405	-0.02713	-0.01802	0.005374	-0.04031
Lag 18	0.000918	-0.00584	-0.00085	0.010818	0.006474	-0.02924	0.016184	-0.01553	-0.02521
Lag 19	-0.00085	0.000445	-0.00187	0.00136	0.000358	-0.01727	-0.02304	0.009723	0.017224
Lag 20	0.004044	-0.00203	0.000765	0.008906	0.009019	0.01095	-0.00304	0.041596	-0.003
Upper Bounds	0.004511	0.00638	0.007813	0.01105	0.014265	0.020174	0.024708	0.034943	0.056344
Lower Bounds	-0.00451	-0.00638	-0.00781	-0.01105	-0.01427	-0.02017	-0.02471	-0.03494	-0.05634

Table 1.3 Autocorrelation of time-aggregated return series

According to Table 1.4, the first lag autocorrelations in time-aggregated returns is significant for almost all the time-aggregated return series. The two exceptions are 15-minute return series and 30-minute return series. Nevertheless, unless the returns are sampled at highest frequencies, the first-lag autocorrelations are very close to zero. This phenomenon is in line with the stylized fact of the absence of linear correlations in return series. Besides, although most of the first-lag autocorrelations are negative, the 78-minute return series has a significantly positive autocorrelation around 0.078.

	T=400	T=800	T=1200	T=2400	T=4000	T=8000	T=12000	T=24000	T=62000
Lag 1	-0.12754	-0.06112	-0.10101	0.005484	0.010487	-0.00232	-0.00546	-0.01872	0.027587
Lag 2	0.002322	-0.00333	0.004357	0.007805	-0.00046	-0.01032	0.000462	-0.00032	-0.01767
Lag 3	-0.00109	0.007604	0.002229	-0.00219	-0.00656	0.008421	-0.00591	0.054872	0.046572
Lag 4	-0.00088	0.002633	0.001115	-0.00251	-0.00617	-0.00937	-0.00105	0.011778	-0.01994
Lag 5	-0.00343	0.003307	0.003341	-0.00115	-0.00343	-0.00122	0.001642	-0.00843	-0.03589
Lag 6	0.002837	-0.00219	-0.00072	-0.00348	0.003181	-0.00626	0.031672	0.018822	0.011799
Lag 7	0.003249	0.001057	-0.0041	-0.00529	0.005893	-0.00176	0.004931	0.01619	-0.00944
Lag 8	0.004238	-0.0015	0.001358	0.002251	-0.0104	0.010009	0.02126	0.002104	0.001322
Lag 9	-0.00333	-0.00418	0.003063	-0.0027	-0.00198	0.006049	-0.00781	-0.00746	0.014582

Lag 10	0.000159	-0.00195	-0.00408	1.21E-05	0.000255	0.023248	-0.01098	-0.01453	-0.00506
Lag 11	0.006263	0.002051	0.001855	0.001977	-0.00095	-0.00385	0.00202	-0.01981	-0.00109
Lag 12	-0.00264	0.000883	0.002637	-0.00164	-0.00241	0.009951	0.015664	-0.00972	0.017795
Lag 13	0.002065	0.003162	-0.0047	-0.00599	0.003079	0.011366	0.01658	0.012838	0.027761
Lag 14	-0.0022	0.004046	-0.00562	0.003244	-0.00645	0.000837	0.002507	-0.02633	-0.04923
Lag 15	0.00328	-0.00585	-0.00197	0.000205	0.003631	-0.00603	0.004895	0.003224	-0.01773
Lag 16	-0.00067	0.000128	0.004372	0.007037	-0.0043	0.016433	0.011076	-0.01657	0.028164
Lag 17	-0.00299	-0.00244	-0.00147	-0.00778	0.011854	-0.01736	-0.00331	-0.00547	-0.03656
Lag 18	-0.00075	0.004705	-0.00096	-0.00153	0.004146	0.015651	-0.01586	0.008453	0.019041
Lag 19	-1.2E-06	-0.00083	-0.00192	0.002899	0.008427	-0.00475	-0.00504	-0.01921	0.062479
Lag 20	-0.00224	-0.00407	0.000604	0.000509	0.010229	0.009002	0.001028	-0.03972	0.017136
Upper Bounds	0.004492	0.006358	0.007791	0.01103	0.014263	0.020241	0.024868	0.035567	0.058596
Lower Bounds	-0.00449	-0.00636	-0.00779	-0.01103	-0.01426	-0.02024	-0.02487	-0.03557	-0.0586

Table 1.4 Autocorrelation of transaction-aggregated return series.

Table 1.5 presents the autocorrelations of the transaction-aggregated return series. There are four return series that do not exhibit significant first lag autocorrelations. Specifically, these return series are the *T-8000*, *T-12000*, *T-24000* and *T-62000* return series.

	30S	1Min	1.5Min	3Min	5Min	10Min	15Min	30Min	78Min
Lag 1	0.491988	0.477804	0.449886	0.382291	0.179437	0.169969	0.19106	0.16986	0.209958
Lag 2	0.006643	0.027125	0.036797	0.042005	0.237255	0.181569	0.2425	0.245731	0.300091
Lag 3	0.00569	0.028009	0.014855	0.054265	0.183269	0.215181	0.197092	0.186489	0.18471
Lag 4	0.006257	0.020629	0.018521	0.045887	0.192326	0.272145	0.206693	0.196917	0.116017
Lag 5	0.007339	0.017769	0.016264	0.040648	0.176035	0.215015	0.158317	0.12894	0.146053
Lag 6	0.006126	0.014964	0.019175	0.043493	0.200129	0.174359	0.167453	0.189599	0.059737
Lag 7	0.005678	0.014957	0.019746	0.040234	0.225508	0.208508	0.147395	0.15568	0.071832
Lag 8	0.004716	0.013338	0.018701	0.046033	0.190529	0.171067	0.134145	0.107617	0.179149
Lag 9	0.004742	0.019658	0.017874	0.04368	0.186128	0.197095	0.147343	0.136254	0.148762
Lag 10	0.00469	0.018044	0.019572	0.038323	0.182269	0.145225	0.137545	0.177981	0.212956
Lag 11	0.003876	0.015657	0.017815	0.055815	0.163151	0.178789	0.13029	0.19438	0.209694
Lag 12	0.003417	0.0162	0.01644	0.044838	0.156464	0.13976	0.103214	0.124865	0.065994
Lag 13	0.004616	0.012558	0.019198	0.038612	0.15703	0.156514	0.12698	0.172143	0.126963
Lag 14	0.004674	0.012609	0.025141	0.036234	0.178204	0.168137	0.142269	0.167073	0.088896
Lag 15	0.003452	0.013802	0.020011	0.036547	0.188024	0.129442	0.11328	0.130477	0.126448
Lag 16	0.004047	0.01249	0.018113	0.036811	0.124292	0.142727	0.125913	0.102856	0.165943
Lag 17	0.004136	0.012437	0.020331	0.037784	0.217291	0.133387	0.121344	0.112751	0.062895
Lag 18	0.0049	0.013744	0.017288	0.036481	0.145389	0.115057	0.123119	0.114057	0.070222
Lag 19	0.005677	0.014089	0.015864	0.041073	0.125791	0.11942	0.175337	0.075653	0.074356
Lag 20	0.004216	0.016683	0.014449	0.0384	0.131843	0.106132	0.124039	0.132469	0.087552
Upper Bounds	0.004511	0.00638	0.007813	0.01105	0.014265	0.020174	0.024708	0.034943	0.056344
Lower Bounds	-0.00451	-0.00638	-0.00781	-0.01105	-0.01427	-0.02017	-0.02471	-0.03494	-0.05634

Table 1.5 Autocorrelation of time-aggregated squared return series

Table 1.6 describes the autocorrelation of time-aggregated squared returns at different frequencies. The volatility clustering phenomenon is obvious. All the time-aggregated returns exhibit strong first lag autocorrelation. These first lag autocorrelation are highly significant and are of unneglectable values. It roughly decays with the decreasing return frequencies.

	T=400	T=800	T=1200	T=2400	T=4000	T=8000	T=12000	T=24000	T=62000
Lag 1	0.499662	0.497506	0.499421	0.000199	0.098132	0.081313	0.068692	0.101727	0.065849
Lag 2	3.05E-05	0.000845	0.000485	0.000241	0.133706	0.103381	0.137646	0.113665	0.104162
Lag 3	3.68E-06	0.00086	0.000298	0.00113	0.096429	0.154742	0.111972	0.123646	0.080449
Lag 4	3.56E-05	0.000816	7.1E-05	0.000495	0.113843	0.119533	0.108117	0.120692	0.116854
Lag 5	6.55E-05	0.000611	7.78E-05	0.000228	0.138599	0.09851	0.109702	0.154277	0.074213
Lag 6	5.99E-05	0.000722	0.000194	0.00031	0.105519	0.098749	0.101987	0.113427	0.051343
Lag 7	3.09E-05	0.001001	0.000192	0.000293	0.135665	0.136247	0.105315	0.095619	0.09716
Lag 8	0.000192	0.000585	-9.7E-06	0.000219	0.0904	0.119155	0.087647	0.121837	0.039991
Lag 9	0.000157	0.000387	-4.2E-05	0.000297	0.085965	0.113032	0.15414	0.121664	-0.01079
Lag 10	-7.8E-06	0.00043	-9.4E-05	0.000115	0.108049	0.103823	0.101166	0.125147	0.061449
Lag 11	-2.6E-05	0.000417	-8.7E-05	0.000107	0.099834	0.112571	0.077766	0.079254	0.020576
Lag 12	-3E-05	0.000309	-6E-05	0.000112	0.094312	0.09446	0.099905	0.042135	0.019021
Lag 13	-1.7E-05	0.00063	0.000241	0.000128	0.093963	0.087801	0.099941	0.108344	0.05221
Lag 14	4.25E-05	0.00079	0.000316	0.000338	0.088012	0.15419	0.094019	0.128008	0.069932
Lag 15	1.89E-05	0.000869	1.62E-05	0.000695	0.103657	0.097573	0.078355	0.093397	0.065214
Lag 16	1.44E-05	0.000618	-5.7E-05	0.001007	0.110921	0.090875	0.104849	0.110485	0.032774
Lag 17	4.33E-05	0.000268	-5.5E-05	0.000353	0.086435	0.086018	0.065852	0.080172	0.020441
Lag 18	2.97E-06	0.000267	0.000161	0.000406	0.108021	0.118912	0.09593	0.103208	-0.00308
Lag 19	3.57E-05	0.000296	0.00016	0.000328	0.087713	0.098797	0.084226	0.101075	0.047338
Lag 20	-1.3E-05	0.000477	-0.0001	5.51E-05	0.091522	0.081316	0.06334	0.075544	-0.00127
Upper Bounds	0.004492	0.006358	0.007791	0.01103	0.014263	0.020241	0.024868	0.035567	0.058596
Lower Bounds	-0.00449	-0.00636	-0.00779	-0.01103	-0.01426	-0.02024	-0.02487	-0.03557	-0.0586

Table 1.6 Autocorrelation of transaction-aggregated squared return series

Table 1.7 exhibit very similar results to Table 1.6. The main difference is that when transaction-aggregated returns are measured at relatively low frequencies the first lag autocorrelation is very small and nearly insignificant. Specifically, the first autocorrelation of $T=62000$ return series is 0.066 that is close to the upper bound of 0.059. Moreover, the first lag autocorrelation is close to zero.

In order to further investigate the difference of the two sampling methods, we discuss and compare several statistical characteristics of returns from the two sampling methods such as aggregational normality and tail properties.

1.3 Aggregational Normality

Although the mean of transaction-aggregated returns and time-aggregated returns can be plausibly assumed to be zero, the excess kurtosis and non-zero skewness clearly suggests that the returns are not normally distributed. Nevertheless, the generally decreasing kurtosis of return series implies that the empirical return distributions are more and more close to normal distributions as the frequencies of returns decrease. This phenomenon is widely identified and discussed as the “aggregational normality” in empirical research using high frequency data (Cont 2001).

The assumption of normally distributed asset returns plays a decisive role in almost all pricing relevant activities such as option valuation, risk management and portfolio construction. Nevertheless, the presence of leptokurtosis and non-zero skewness is a widely investigated and accepted stylized fact regarding the distribution of asset returns (Cont 2001). Fama (1965) showed that the distribution of returns approached the normal distribution if the holding period extends from one day to one month. Mandelbrot (1963) used a stable process to explain the deviation of equity returns from the normal distribution. Ane and Geman (2000) showed that the high frequency returns are conditional on the number of transactions. Moreover, they argued that, based on Clark’s work on 1973, the valuation of an equity is equivalent to the identification of the time deformation of original high frequency return series that could recover the normality of such series.

Motivated by these results and the fact that time-aggregated returns and transaction-aggregated returns have different statistical characteristics such as kurtosis, I compare the two sampling methods in terms of the aggregational normality. I want to address an empirical question: “Which sampling method can provide normally distributed return series with higher frequency?”. The answer to this question might help us to understand

the nature of tick-by-tick returns and further contribute to its dynamic modelling.

I carry out two sets of normality tests. In the first set of normality test, I examine the normality of the transaction-aggregated return series and time-aggregated return series using the whole sample. In the second set of normality test, I examine the normality of the two sets of return series on daily basis. Specifically, I carry out the normality test on each day's transaction-aggregated return series and time-aggregated return series. The statistics used in the normality test are EDF statistics D , V , W^2 , U^2 and A^2 . The details of these statistics and the test procedures are presented in Appendix A (Stephens 1974).

	30Sec	1Min	1.5Min	3Min	5Min	10Min	15Min	30Min	78Min
Sample size	196560	98280	65520	32760	19656	9828	6552	3276	1260
D	44.38997	28.18536	22.82909	15.0778	10.53782	7.185938	5.868893	4.22834	2.763899
V	88.36515	55.86088	45.43276	29.79012	20.05774	13.8916	11.09009	7.852333	5.134825
W^2	975.4515	354.4086	237.9483	104.0238	45.52168	21.96869	13.26837	6.698358	2.660654
U^2	975.4213	354.3744	237.9338	103.9936	45.47667	21.92438	13.18956	6.668071	2.635175
A^2	Inf	Inf	Inf	Inf	Inf	Inf	75.5676	37.0617	14.71987

Table 1.7 EDF statistics of time-aggregated returns

	T=400	T=800	T=1200	T=2400	T=4000	T=8000	T=12000	T=24000	T=62000
Sample size	198261	98942	65893	32876	19663	9763	6468	3162	1165
D	31.70505	13.65617	12.30105	5.029308	2.455067	1.334636	1.37649	0.912603	0.666228
V	62.59288	27.18621	24.54178	9.948603	4.589417	2.466574	2.405428	1.995015	1.296045
W^2	339.6452	59.89948	62.45164	10.39527	1.562107	0.382361	0.334116	0.300721	0.081262
U^2	339.6411	59.89944	62.45128	10.3944	1.562038	0.372619	0.31038	0.275914	0.075185
A^2	Inf	Inf	Inf	Inf	9.000769	2.763163	2.265547	1.75191	0.51716

Table 1.8 EDF statistics of transaction-aggregated returns

Statistics	Significance level		
	5%	2.5%	1%
D	0.895	0.955	1.035
V	1.489	1.585	1.693
W^2	0.126	0.148	0.178
U^2	0.116	0.136	0.163
A^2	0.787	0.918	1.092

Table 1.9 Critical values of EDF statistics for the null hypothesis of normal distribution

I conclude two important results from Table 1.8. First, the values of the EDF statistics decrease monotonically as the frequencies of the time-aggregated returns decrease. For instance, the values of the Kolmogorov statistics D , Kuiper statistic V and the Cramér–von

Mises statistics W^2 drop from 44.390, 88.365 and 975.425 for the 30-second return series to 2.764, 5.135 and 2.661 for the 78-minute return series respectively. The phenomena that the values of these EDF statistics approach to the critical values in Table 1.10 suggest that the distributions of time-aggregated return series approach to the normal distribution as the return frequencies decrease. This is in line with the widely identified stylized fact of aggregational normality.

Second, the null hypothesis is rejected for all time-aggregated return series at significance level of 1% using all these five EDF statistics. For instance, the values of the D , V , W^2 , U^2 and A^2 for the 78-minute return series are 2.764, 5.135, 2.661, 2.634 and 14.720 respectively and meanwhile the critical values of these statistics are 1.035, 1.683, 0.178, 0.163 and 1.092 respectively at significance level of 1%. This empirical finding suggests that the empirical distributions of time-aggregated returns are not normally distributed although they approach to the normal distribution as the return frequencies decrease.

With respect to the results for the transaction-aggregated return series in Table 1.9, the situation is very similar. The aggregational normality is identified by the monotonical decreasing values of the EDF statistics. Recall that since I set the values of T that is the fixed number used to aggregate the tick-by-tick observations according to sample size of time-aggregated return series, I can compare their statistics accordingly in this sense. When compared with Table 1.8, Table 1.9 shows two interesting findings:

First, the values of the statistics of the transaction-aggregated return series are much smaller than the values of these statistics of corresponding time-aggregated return series. E.g. The values of statistics D , V , W^2 and U^2 for the return series aggregated from every 400 transactions are 31.705, 62.592, 339.645 and 339.641 respectively while the values of these statistics for the 30-second return series are 44.390, 88.365, 975.451 and 975.421 respectively. It implies that the empirical distributions of transaction-aggregated returns are more like the normal distribution when compared with the time-aggregated returns of similar sample size.

Second, the values of statistics D , V , W^2 , U^2 and A^2 for the $T-62000$ return series are 0.666, 1.296, 0.081, 0.075 and 0.517 respectively. The null hypothesis that the sample comes from a normal distribution cannot be rejected at significance level of 5% when using any of these statistics. I can therefore assume the returns aggregated from every 62000 tick-by-tick transactions following the normal distribution. In contrast, the values of statistics D , V , W^2 , U^2 and A^2 of the corresponding 78-minute return series all obviously exceed that critical values at significance level of 1%. Meanwhile, the value of Kolmogorov statistics D of the $T-24000$ return series is 0.913 that is less than the critical value of D at significance level of 2.50%.

These results suggest an important empirical difference between the two sampling methods. Compared with returns sampled from fixed time interval, returns sampled from fixed number of transactions are more like a population from normal distribution. Moreover, when the return frequency decreases, the empirical distribution of transaction-aggregated returns approaches the normal distribution faster than the time-aggregated returns.

For a symbol with high liquidity like SPY, there are hundreds of thousands tick-by-tick transactions during each day. Thus, I can further investigate this difference by examining the aggregation normality of the two sampling methods on daily basis. The second set of tests is for the further illustration and completeness. Specifically, I carry out the normality test using EDF statistics D , W^2 and A^2 on each day's time-aggregated and transaction aggregated returns. As showed by Stephens (1974), in the case of testing unknown normal distributions, W^2 and A^2 are preferred statistics since they have stronger power than D . D is chosen for the completeness. Returns within a trading day are considered to follow normal distribution if all the three statistics D , W^2 and A^2 cannot reject the null hypothesis at significance level of 5%.

Rate of normality is defined to be the ratio of the number of days that all three EDF tests cannot reject the null hypothesis to 252 that is the total number of the days in our data.

In order to allow for enough observations, I use $\Delta t=0.5, 1, 1.5, 3, 5$ minutes for time-

aggregated returns within each trading day. The corresponding transaction-aggregated returns within each trading day are sampled from $T=400, 800, 1200, 2400$ and 4000 . Note that, the sample size of transaction-aggregated transactions on each day varies largely and meanwhile the sample size of time-aggregated returns on each day is fixed. The least number of transactions occurring in one trading day is 121489. The $T=4000$ return series within that trading day has 30 observations.

	30S	1Min	1.5Min	3Min	5Min
Sample size of each day	780	390	260	130	78
Rate of normality	1.59%	9.52%	14.29%	36.51%	59.13%

Table 1.10 Summary of within-day normality test on time-aggregated returns

	T=400	T=800	T=1200	T=2400	T=4000
Average sample size of each day	785	391	261	130	78
Rate of normality	3.17%	30.56%	84.92%	89.92%	92.21%

Table 1.11 Summary of within-day normality test on transaction-aggregated returns

Table 1.11 indicates that the rate of normality increases monotonically as the length of the fixed time interval used to sample returns for the time-aggregated return series. This results further support the aggregational normality. Specifically, it increases from 1.59% for the 30-second within-day return series to 59.13% for the 5-minute within-day return series. It implies that when the return series during a trading day are sampled from every 5 minute, there are more than 40% of them cannot be assumed to follow normal distributions.

Table 1.12 suggests that the rate of normality for transaction-aggregated returns also increases monotonically as the number of transactions used to calculate returns increases. However, an important difference is that transaction-aggregated returns in each day have much higher probability to follow normal distribution when compared with time-aggregated returns. For instance, when we aggregated every 1200 tick-by-tick transactions to generate the return series in each of the 252 trading days, 84.92% of these return series can be considered as normally distributed. Meanwhile, the average sample size of the return series aggregated from 1200 tick-by-tick transactions in each day is around 260. In contrast, only 14.29% of the 1.5-minute return series in each day can be assumed to follow normal distribution. More obvious evidence comes from the last column of Table 1.12, when I aggregate every 4000 tick-by-tick transactions to generate return series in each day, the average sample size of each trading day is around 78 while the normality rate is 92.21%.

Meanwhile, the corresponding 5-minute return series within each day have a normality rate of 59.13%.

The test of the return series sampled from fixed time interval and fixed number of transactions with respect to the aggregational normality suggests that the empirical distributions of transaction-aggregated returns are closer to the normal distribution when compared with the empirical distribution of time-aggregated returns. Moreover, compared with sampling from fixed time interval, sampling from fixed number of transactions can recover the normality of return series at lower costs in the sense it can preserve larger sample size.

The aggregational normality is not surprising at first glance due to the Central Limit Theorem. Recall the standard Central Limit Theorem : *Let $\{X_t\}$ be a sequence of identically distributed independent random variables with $E(X_t) = u$ and $Var(X_t) = \sigma^2$ for all t . Then,*

$\lim_{t \rightarrow \infty} P\left(\frac{\sum_{t=1}^n X_t - nu}{\sqrt{no}} \leq x\right) = \Phi(x)$, where Φ is the cumulative distribution function of the standard normal distribution.

i.i.d. assumption is not valid for the $\{X_t\}$ of tick-by-tick data. Jirak (2016) and Hormann (2009) discuss some generalizations of the standard Central Limit Theorem for economic time series. Suppose the tick-by-tick return process $\{X_t\}$ is modeled by a $GARCH(p,q)$ process

$$X_t = \varepsilon_t \sigma_t, \varepsilon_t \sim N(0,1) \text{ and } \{\varepsilon_t\} \text{ is an i.i.d sequence,}$$

$$\sigma_t^2 = \mu + \sum_{i=1}^p \alpha_i \sigma_{t-i}^2 + \sum_{j=1}^q \beta_j X_{t-j}^2 \text{ where } \alpha_i \text{ and } \beta_j \text{ are coefficients.}$$

A stochastic process is said to be stationary if it has an unconditional joint distribution that does not change when shifted in time. For detailed discussion regarding stationarity of GARCH process, see e.g., DeGroot and Schervish (2014) and Walter (2010). Suppose that the stationarity of $\{X_t\}$ can be verified under the $GARCH(p,q)$ specification, then the Central Limit Theorem still holds for $\{X_t\}$. The stationarity is determined by the parameters α_i 's and β_j 's. Specifically, the process is said to have weak stationarity if $\sum_{i=1}^p \alpha_i +$

$\sum_{j=1}^q \beta_j < 1$. However, the tick-by-tick return data presents many erratic statistical features that cannot be explained by the GARCH models. The GARCH models are also widely reported as not suitable for modelling high frequency return series (Engle 2000 and Hasbrouck 1996). Moreover, the assumed stationarity for the GARCH process generally is not valid when using high frequency data. A more detailed discussion regarding Berry-Essen bound for GARCH process is presented in Appendix B.

An alternative explanation of the aggregational normality is provided by the mixture of distributions model (MODM). It assumes the variability of equity prices are caused by the differences in information arrival rates. The standard MODM model assumes traders with different expectations, different risk profiles and corresponding reservation prices (Harris 1986 and Richardson and Smith 1994). Traders adjust their reservation prices as the information arrives and the market price is moved consequently. Thus, the daily price change is the sum over the within-day price changes. Then Central Limit Theorem shows that the daily price changes can be described as the mixtures of independent normal distributions.

From the perspective of MODM framework, if we view the transactions are results of new information arrival, then returns aggregated from fixed number transactions are proxies of the sum of the unseen information flows. Therefore, the Central Limit Theorem can provide a plausible explanation of the aggregational normality. However, the effects of information on traders and trading are not addressed in this framework. Besides, for my tick-by-tick data transaction records, the average waiting-time between two consecutive transactions is around 0.07s. It is unplausible to assume that the traders can respond to new information arrivals in such a short time interval.

1.4 Tails of the return distributions

My empirical analyses suggest that intraday return series, sampled from fixed time intervals and fixed number of transactions, both exhibit heavy tails. This non-Gaussian character of the high frequency return distribution implies the non-negligible probability of occurrence of abnormal market price movement. As a result, an accurate description of

the tails of return distribution is desirable. The use of generalized Pareto distribution for measuring distribution tails is introduced by DuMouchel (1983). After that, it is widely accepted in empirical financial research for the modelling of the tails of return distributions (Jasen and de Vires 1991, Loretan and Phillips 1994, and Young and Graff 1995). The interpretation of the tail thickness is often related with the tail index α of stable distributions. In this section, I first discuss literatures of measuring returns tail thickness. Then I analyze some basic properties of the stable distributions and generalized Pareto distribution. Finally, I present my tail fit using generalized Pareto distribution and compare the empirical results from the different sampling methods.

Mantegna and Stanley (1995) argued that the distribution of high-frequency S&P index returns can be well described by a truncated Lévy distribution. The density of Lévy distribution has an approximately power-law decay in tail. Specifically, $S(r) \sim r^{-\alpha}$, where $S(r) = 1 - \Pr(R \leq r)$ and $0 < \alpha < 3$. Research focusing on the asymptotic behavior of the tails of high-frequency returns in the US and the Chinese stock markets suggest that the high frequency returns tend to follow an inverse cubic law in the tail where the power-law exponent α is found to be close to 3 (Gopikrishnan et al. 1998, Gu et al. 2007 and Plerou et al. 1999). Empirical research on this topic for other markets can be found in Makowiec and Gnaciński (2001), Bertram (2004) and Coronel-Brizio et al. (2005). It is also noticed by researchers that the tail distribution of returns evolves from power law to Gaussian as the return frequencies decrease (Ghashghaie et al. 1996 and Castaing et al. 1990).

The class of stable distributions is introduced by Paul Lévy in his research regarding the sum of independent identically distributed random variables in the 1920's. According to the generalized Central Limit Theorem, if the sum of a sequence of independent identically distributed random variables has a limiting distribution, the limiting distribution must belong to the class of stable distributions. As a result, it is natural to assume that the asset returns should approximately follow a stable distribution. The use of stable distributions to describe the heavy tails and skewness of asset returns can be found in Mandelbrot (1963), Fama (1965) and Rachev and Mitnik (2000).

Generally, a random variable X is said to be stable provided that any positive linear combinations of independent copies of X follow distributions that are the same as that of X . A rigorous mathematical presentation of the definition of stable distribution can be found in Nolan (2018). As argued by DuMouchel (1975), stable distributions can be defined concisely by the log characteristic function. Specifically,

$$\log E(\exp(iXt)) = \begin{cases} i\delta t - |ct|^\alpha \left[1 - i\beta \operatorname{sgn}(t) \tan \frac{\pi\alpha}{2} \right], & \alpha \neq 1; \\ i\delta t - |ct| \left[1 + i\beta \frac{2}{\pi} \operatorname{sgn}(t) \log |ct| \right], & \alpha = 1. \end{cases}$$

The shape of the stable distributions is governed by

- the characteristic exponent $\alpha \in (0, 2]$,
- the skewness parameter $\beta \in [-1, 1]$,
- the location parameter $\delta \in R$, where R is the set of real numbers, shifts the distributions along the real line,
- $c \in (0, \infty)$ is the scale parameter which expands or contracts the distribution around δ .

In the case of $\alpha = 2$, the distribution reduces to a normal distribution with variance $2c^2$ and mean δ .

A basic property of the stable distributions is, when $\alpha < 2$ the distributions have Pareto tails. Suppose that X is a positive random variable. The survival functions $S(X) = 1 - \Pr(X \leq x)$ behave asymptotically like $x^{-\alpha}$ for large x and yields infinite absolute moments of order higher and equal to α . Thus, the variance is infinite.

In his influential paper in 1983, DuMouchel introduced the use of generalized Pareto distribution for modelling distribution tails. The generalized distribution can be characterized as:

$$f(x|k, \sigma, x_0) = \left\{ \begin{array}{ll} \frac{1}{\sigma} \left(1 + \frac{k(x - x_0)}{\sigma} \right)^{-1-1/k} & ; x > x_0 \text{ and } k > 0 \\ \frac{1}{\sigma} \exp \left(-\frac{x - x_0}{\sigma} \right) & ; x > x_0 \text{ and } k = 0 \\ \frac{1}{\sigma} \left(1 + \frac{k(x - x_0)}{\sigma} \right)^{-1-1/k} & ; x_0 < x < x_0 - \frac{\sigma}{k} \text{ and } k < 0 \end{array} \right\}$$

where the $k \in R$ is the shape parameter and $\sigma \in (0, \infty)$ is the scale parameter. x_0 is the threshold parameter.

In the case $k > 0$, the upper tail probabilities behave asymptotically like $x^{-\alpha}$ where $\alpha = 1/k$. The Pareto distribution with exponent parameter $1/k$ can be viewed as a special case of the generalized Pareto distribution with $k > 0$ and $\sigma = kx_0$. The most important advantage of the generalized Pareto distribution in the measurement of thickness of tails is that it allows a continuous range of possible shapes that include exponential, Pareto and normal distributions. As a result, it allows the tails to “speak for themselves” as possibly as they can.

A common method of estimating the tail index of the unknown stable distribution of returns related with generalized Pareto distribution is organized as follows: First, one can use extreme observations to fit the generalized Pareto distribution and estimate k by maximum likelihood estimation. The maximum likelihood estimators are provided by DuMouchel (1983). Second, one can further use $1/k$ to estimate the tail index α .

The generalized Pareto distribution allows the $\alpha > 2$. In this case, returns are usually considered to reassuringly have finite variance. However, as argued by McCulloch (1997), there are stable distributions with α smaller than 2 for which the tail index estimated from the generalized Pareto distribution is larger than 2.

As a result, if the true distribution of returns is stable, measuring tail thickness cannot be a reliable method of estimating the stable index α . He therefore further argues that it is necessary to use the whole sample instead of just tails to estimate the tail index for stable distributions.

Based on previous discussion, instead of imposing stable assumptions on returns, I use the positive and negative extreme observations to study the tail behaviors of the return distribution. We use the generalized Pareto distribution to model the tails of transaction-aggregated returns and time-aggregated returns.

	Positive Tails				Negative Tails			
	k	Standard Error	σ	Standard Error	k	Standard Error	σ	Standard Error
30S	0.213349	0.010734082	0.00014	2.06738E-06	0.212478942	0.010842113	0.000142	2.0803E-06
1Min	0.191014	0.015283361	0.00019	4.11984E-06	0.194646812	0.01412995	0.000212	4.23016E-06
1.5Min	0.189211	0.018536001	0.00024	6.14442E-06	0.18108007	0.018399117	0.000245	6.1647E-06
3Min	0.128798	0.028409838	0.00033	1.23044E-05	0.209424795	0.023131603	0.000365	1.22908E-05
5Min	0.143225	0.036490234	0.00043	2.07638E-05	0.008807853	0.031675092	0.000499	2.24379E-05
10Min	0.105634	0.049591498	0.00061	4.08831E-05	0.071263276	0.050239157	0.000635	4.28631E-05
15Min	0.154077	0.067075508	0.00069	5.96246E-05	0.02000227	0.060740059	0.000820	6.73529E-05
30Min	0.073589	0.084268123	0.00099	0.000114679	0.10502430	0.091534969	0.000921	0.000110624
78Min	0.087694	0.152099476	0.00166	0.000328977	-0.02421194	0.121891616	0.002052	0.000359752

Table 1.12 The modelling of tails of time-aggregated returns using generalized Pareto distribution

	Positive Tails				Negative Tails			
	k	Standard Error	σ	Standard Error	k	Standard Error	σ	Standard Error
T=400	0.110184	0.007648207	9.66E-05	1.21462E-06	0.105233	0.00743404	9.84E-05	1.22563E-06
T=800	0.104141	0.010710156	0.000134	2.37062E-06	0.109946	0.011215493	0.000132	2.38042E-06
T=1200	0.111675	0.013679815	0.000163	3.59781E-06	0.099391	0.012689702	0.000161	3.45268E-06
T=2400	0.107225	0.01876058	0.000216	6.66775E-06	-0.01781	0.022911138	0.000238	8.00696E-06
T=4000	-0.01577	0.02886951	0.000327	1.40729E-05	-0.02668	0.025747866	0.00031	1.27124E-05
T=8000	0.023837	0.041332527	0.000431	2.64135E-05	-0.12921	0.044558396	0.000481	3.04163E-05
T=12000	-0.01089	0.046235937	0.000535	3.86849E-05	-0.16414	0.050360159	0.000602	4.48067E-05
T=24000	0.097219	0.082951777	0.000639	7.33057E-05	-0.02384	0.093969857	0.000643	7.91183E-05
T=62000	-0.29957	0.131460548	0.001585	0.000287822	-0.13987	0.123711546	0.001118	0.00020057

Table 1.13 The modelling of tails of transaction-aggregated returns using generalized Pareto distribution

I use the lower 5% and upper 5% of each return series to define the negative tails and positive tails respectively. Estimation of parameters are based on the maximum likelihood estimation.

For the positive tail of time-aggregated returns, the shape parameter k decreases almost monotonically with decreasing frequencies of returns. Specifically, it drops from 0.213 for the 30-second return series to 0.088 for 78-minute return series. Recall that when $k=0$, the

generalized pareto distribution reduces to the exponential distribution. Consequently, sample from distributions whose tails decay exponentially or asymptotically exponentially should give values of k close to 0 in the estimation. In contrast, positive values of k suggest that the tails can be roughly interpreted as the tail of a Pareto distribution with shape parameter $\alpha = \frac{1}{k}$.

The negative tail of time-aggregated returns behaves differently. The shape parameter k bounce erratically between 0.212 for the 30-second return series and 0.105 for 30-minute return series across the first eight arrays in Table 1.9. For the negative tails of 30-second returns, 1-minute return and 1.5-minute returns, the values of k are 0.212, 0.154 and 0.182 respectively, which suggests thicker tails compared with normal distribution. The shape parameter k of the negative tail of 78-minute return series is negative. It can be interpreted as a thinner tail behavior compared with the exponential distribution.

With respect to the positive tail of transaction-aggregated returns, the shape parameter varies from -0.300 for returns aggregated by every 62000 tick-by-tick transactions to 0.110 for returns aggregated by every 400 tick-by-tick transactions. There is no obvious monotonicity in the values of k against the frequencies of returns.

The negative tails of transaction-aggregated returns exhibit interesting characteristics. The values of shape parameter k decrease almost monotonically as the decreasing frequencies of returns. It starts with a value of 0.105 for the $T-400$ return series, suggesting thicker negative tail compared with normal distribution. It approaches 0 for the $T-2400$ return series and $T-4000$ return series, indicating that the negative tail of the return distributions is thinner than the tails of exponential distribution.

The values of σ estimated for both positive and negative tails of time-aggregated returns with different frequencies are very close to 0. Besides, the values of k for positive tail and negative tail of the same return series are obviously different from each other, suggesting the return distribution should be asymmetric. In comparison, time-aggregated returns have positive tails that approaches the tail of normal distribution as the return frequencies

decrease while transaction-aggregated returns have negative tail that approaches the tail of normal distribution with the decreasing return frequencies.




Figure 1.3 Positive tail of SPY 5-minute returns




Figure 1.4 Negative tail of SPY 5-minute returns

Figures 1.3 and 1.4 give a graphic illustration for the use of generalized Pareto distribution to fit the tails of high frequency SPY returns.

1.5 Realized volatility and standardized returns

In this section, I compare the variabilities of the time-aggregated returns and transaction-aggregated returns. Return volatilities plays a decisive role in theoretical development of asset pricing, and risk management. Due to the extremely high liquidity of SPY, market microstructural issues such as “price discreteness”, “bid-ask” bounces, “split transactions”,

recording accuracy and intraday return seasonality have very strong impact on the pattern of tick-by-tick returns. A direct parametric modelling such as GARCH-ARCH and stochastic volatility models therefore might be hazardous. A widely accepted nonparametric volatility measurement of high frequency return is the Realized Volatility (RV). RV measures the volatility for a time period by the sum of squared tick-by-tick returns during that period. It is based on the following reasonings.

Suppose that the asset prices S_t follows the continuous time diffusion of

$$\begin{aligned} X_t &= \ln(S_t), \\ dX_t &= u_t dt + \sigma_t dW_t, \end{aligned} \quad (1.1)$$

where W_t is a standard Brownian motion, u_t is the drift coefficient, σ_t^2 is the instantaneous variance that should be squared integrable, i.e., $E(\int_0^t \sigma_s^2 ds) < \infty$. In empirical applications using high frequency data, the u_t is commonly assumed to be a constant zero and σ_t is modelled by stochastic volatility process such as the Cox-Inger-Ross process applied in the Heston model (Heston 1993). Thus, one can immediately infer from the stochastic equation (1.1) that

$$X_t - X_0 = \int_0^t \sigma_s dW_s. \quad (1.2)$$

Integrated variance $IV(t, 0)$ that is the diffusive sample path variation is $IV(t, 0) = \int_0^t \sigma_s^2 ds$. For a stochastic real-valued process X_t , $t \in [0, \infty)$, from a probability space of (Ω, F, P) , the quadratic variation of X_t , $[X_t]$ is defined as the process,

$$[X_t] = \lim_{\text{gap } P \rightarrow 0} \sum_{i=1}^n (X_{t_i} - X_{t_{i-1}})^2,$$

if such limit exists. $P = \{t_0, \dots, t_n\}$ is a partition of the interval $[0, t]$. The convergence is in the sense of convergence in probability measure. In the case of (1.1), since the standard Brownian motion has quadratic variation $[W_t] = t$ almost surely for all $t \in [0, \infty)$, $[X_t] = \int_0^t \sigma_s^2 ds = IV(t, 0)$. The two coincide. An important implication immediately is that for a sequence of regular partitions of $[0, t]$, $\{P_n\}$ with $\lim_{n \rightarrow \infty} \text{gap } P_n \rightarrow 0$,

$$\lim_{n \rightarrow \infty} \sum_{i=1}^n (X_{t_i} - X_{t_{i-1}})^2 = \int_0^t \sigma_s^2 ds, \quad (1.3)$$

where the convergence in (1.3) is in the sense of convergence in probability.

Equation (1.3) justifies the use of sum of squared tick-by-tick log returns to estimate the variance. In other words, the estimation error of the realized volatility dies out with the increasing number of observations. Consequently, the realized volatility computed from the tick-by-tick data should give the most accurate estimation. The realized volatility and its estimation accuracy are widely explored by researchers such as Zhang et al. (2005), Gallant, Hsu, and Tauchen (1999), Chernov and Ghysels (2000), Andersen et al. (2001), Mykland and Zhang (2002) and Barndorff-Nielsen and Shephard (2002).

A main advantage of the realized volatility is that it takes, at the highest frequency, each tick-by-tick observation into account. For instance, when estimating the daily volatility, the intraday price can fluctuate dramatically over the trading day but ends with a value that is very close to the opening price. Thus, volatility measures that are based on daily observations cannot accurately capture the variability of intraday prices. In contrast, the realized volatility gives a more realistic description.

I now present the realized volatility of time-aggregated returns and transaction-aggregated returns. For the transaction-aggregated returns, recall that the i th return on day d of the transaction-aggregated return series is defined as

$$(T)r_{i,d} = \sum_{j=(T-1)(i-1)+1}^{(T-1)i} \ln \left(\frac{p_{j+1,d}}{p_{j,d}} \right) = \ln \left(\frac{p_{(T-1)i+1,d}}{p_{(T-1)(i-1)+1,d}} \right),$$

where T denotes the number of tick-by-tick transactions used to aggregate the returns. According to equation (1.3), the realized volatility $(T)v_{i,d}$ is defined as

$$(T)v_{i,d}^2 = \ln^2 \left(\frac{p_{(T-1)i+1,d}}{p_{(T-1)(i-1)+1,d}} \right).$$

With respect to the time-aggregated returns, the realized volatility is the sum of squared tick-by-tick returns during each time interval.

	$v_{i,d}^2$				$\ln v_{i,d}$			
	Mean	Std	Skewness	Kurtosis	Mean	Std	Skewness	Kurtosis
30S	2.93E-06	0.00045948	439.3788	193962.296	-7.98032	0.658878	1.499165	9.539247
1Min	5.88E-06	0.00064967	310.6506	96988.9004	-7.55913	0.671298	1.705153	9.912378
1.5Min	8.81E-06	0.00079554	253.6311	64666.3046	-7.31379	0.688522	1.714867	9.578119
3Min	1.76E-05	0.00112542	179.0169	32261.9401	-6.89009	0.72649	1.584125	8.573127

5Min	2.94E-05	0.00145411	138.4496	19319.6211	-6.57163	0.752474	1.512823	6.87682
10Min	5.88E-05	0.00206903	97.41162	9595.46353	-6.12909	0.785616	1.241688	5.544807
15Min	8.83E-05	0.00255543	79.32401	6373.98438	-5.8684	0.795599	1.127256	4.88841
30Min	0.000176	0.00368225	55.45966	3136.82158	-5.42233	0.791258	0.955727	4.195112
78Min	0.000459	0.00604078	33.707	1172.51807	-4.82086	0.755883	0.869512	4.159424

Table 1.14 The summary of the realized volatilities of time-aggregated returns

Table 1.15 summaries the descriptive statistics of the realized variance and logarithmic realized volatility of time-aggregated returns. It is very clear that the realized volatilities of time-aggregated returns are right-skewed and leptokurtic. The kurtosis and skewness decrease monotonically with return frequencies.

	$v_{i,d}^2$				$lnv_{i,d}$			
	Mean	Std	Skewness	Kurtosis	Mean	Std	Skewness	Kurtosis
T=400	2.92E-06	0.000324	320.7945	105830.5	-7.75039	0.458618	3.772327	26.27352
T=800	5.84E-06	0.000641	312.2953	97993.27	-7.34845	0.493896	3.717241	22.18606
T=1200	8.77E-06	0.000657	236.5604	58188.26	-7.10984	0.521005	3.528754	19.19153
T=2400	1.49E-05	0.000648	177.2084	31878.44	-6.69446	0.570655	3.040223	13.85339
T=4000	1.9E-05	0.000107	13.50224	276.3899	-6.37836	0.60716	2.61194	10.32108
T=8000	3.82E-05	0.00017	14.01847	371.3236	-5.93413	0.650256	2.049934	6.891467
T=12000	5.74E-05	0.000219	10.48758	181.9913	-5.67099	0.664962	1.778826	5.611981
T=24000	0.000116	0.000389	12.07596	260.4523	-5.21996	0.669629	1.400238	4.237923
T=62000	0.000305	0.000807	7.198153	81.4959	-4.62929	0.647569	1.126266	3.543235

Table 1.15 The summary of the realized volatilities of transaction-aggregated returns

In general, the realized volatilities of transaction-aggregated returns exhibit characteristics similar to the time-aggregated returns. There are two interesting empirical findings that can be concluded from Tables 1.15 and 1.16. First, compared with the time-aggregated returns, the realized variances of transaction-aggregated returns are of much smaller kurtosis and skewness. Second, the means of realized variance of transaction-aggregated returns are constantly lower than time-aggregated returns (the standard deviations of the realized variance of transaction-aggregated returns and time-aggregated returns are equally telling). It implies that the variability of intraday returns measured from fixed time interval is higher than the variability of intraday returns measured from fixed number of transactions. This implication is in line with the microstructural theory of MODM discussed previously in section 1.3.

I now standardize the time-aggregated returns and transaction-aggregated returns by the corresponding realized volatility. The following Tables 1.17 and 1.18 summarize the descriptive statistics of the standardized time-aggregated returns and standardized transaction-aggregated returns respectively.

	$r_{i,d}/v_{i,d}$					
	Mean	Std	Skewness	Kurtosis	Maximum	Minimum
30S	0.003849	0.57724	13.43105	2209.772	70.71066	-44.7077
1Min	0.001744	0.721509	-89.9912	19575.73	58.02649	-149.505
1.5Min	0.005974	0.636411	69.30217	10242.28	99.50478	-22.1161
3Min	0.007103	0.689116	91.77533	13294.85	99.50478	-12.9074
5Min	0.004716	0.395679	1.950666	79.13874	13.7076	-6.03503
10Min	0.006182	0.359443	-0.41667	10.63182	2.18529	-5.85823
15Min	0.009758	0.339246	-0.05126	3.729002	1.802042	-1.38976
30Min	0.011492	0.316745	-0.04838	4.039979	1.689292	-1.27714
78Min	0.018069	0.28456	-0.02162	4.035005	1.049404	-1.12318

Table 1.16 Descriptive statistics of standardized time-aggregated returns at different frequencies

	$r_{i,d}/v_{i,d}$					
	Mean	Std	Skewness	Kurtosis	Maximum	Minimum
T=400	0.000839	0.463358	-0.0016	3.510289	2.494489	-2.51162
T=800	0.001206	0.446944	0.014925	3.407308	2.133432	-2.24665
T=1200	0.001459	0.437086	0.008439	3.395956	2.324799	-1.89998
T=2400	0.00142	0.418892	-0.01437	3.305479	2.22261	-1.85995
T=4000	0.002394	0.405768	-0.00457	3.397714	1.774267	-1.88395
T=8000	0.002631	0.38251	-0.01341	3.286298	1.458433	-1.3815
T=12000	0.003402	0.365811	-0.0599	3.377889	1.456133	-1.36493
T=24000	0.004396	0.33959	-0.12424	3.558695	1.40428	-1.4491
T=62000	0.004716	0.31532	-0.00783	3.57636	1.159459	-1.09435

Table 1.17 Descriptive statistics of standardized transaction-aggregated returns at different frequencies

Compared with the standardized time-aggregated returns, the standardized transaction-aggregated returns are strikingly much more regular. First, in general the skewness of the standardized transaction-aggregated returns is very close to zero. In contrast, the skewness of standardized time-aggregated varies dramatically at high frequencies.

Second, the kurtosis of standardized transaction-aggregated returns fluctuates around 3. Compared with Table 1.3, it is noticeable that the kurtosis of transaction-aggregated returns

at high frequencies significantly reduces. For instance, the kurtosis of raw T -400, T -800 and T -1200 returns are 2878.43, 401.41 and 1046.23 respectively. Meanwhile, the kurtosis of standardized T -400, T -800 and T -1200 returns are 3.51, 3.41 and 3.40 respectively.

These empirical findings suggest that the empirical distributions of standardized transaction-aggregated returns are very close to the normal distribution. I now use EDF statistics to formally test the unconditional distribution of standardized transaction-aggregated returns and time-aggregated returns. As in section 1.3, I examine the normality of intraday standardized returns within each trading day. Intraday returns within a trading day are considered to follow normal distribution if all the three statistics D , W^2 and A^2 cannot reject the null hypothesis at significance level of 5%.

Standardized Transaction-aggregated returns			Standardized Transaction-aggregated returns		
	N_0	$N_0/252$		N_0	$N_0/252$
30S	43	17.06%	T=400	93	36.90%
1Min	137	54.37%	T=800	184	73.02%
1.5Min	165	65.48%	T=1200	212	84.13%
3Min	203	80.56%	T=2400	222	88.10%
5Min	207	82.14%	T=4000	227	90.08%

Table 1.18 Summary of within-day normality test on standardized returns

Table 1.19 present the normality rate of standardized transaction-aggregated returns and time-aggregated returns. The N_0 denotes the number of days whose intraday returns are normally distributed. Compared with Tables 1.11 and 1.12, a salient difference is that when measuring at the highest frequencies (like every 30 seconds for time-aggregated returns or every 400 tick-by-tick transactions for transaction-aggregated returns), the standardized returns are much more higher probabilities of being normally distributed.

Moreover, the standardized transaction-aggregated returns still have higher normality rate than the standardized time-aggregated returns. It further confirms my previous empirical finding that the transaction-aggregated returns are more likely of being normally distributed.

I end this chapter by a brief discussion of the daily SPY returns. The realized variance of a

trading day is calculated by the sum of squared tick-by-tick returns.

	Mean	Std	Skewness	Kurtosis	Min	Max	A^2	Critical values
Raw daily returns	0.0001	0.006	-0.6354	4.744	0.0172	-0.0215	3.2636	0.7496
Standardized daily returns	0.0326	0.2152	0.0287	3.2914	0.7575	-0.6726	0.3088	0.7496

Table 1.19 Summary of SPY daily returns

Table 1.20 suggests that the raw SPY daily returns are left-skewed and leptokurtic. In contrast, the standardized daily SPY returns have skewness very close to zero and only are slightly leptokurtic. Moreover, as suggested by the last two columns in Table 1.20, the Andersen-Darling statistics cannot reject the null of normal distribution at significance level of 5% for the standardized daily returns.

	v_i^2				lnv_i			
	Mean	Std	Skewness	Kurtosis	Mean	Std	Skewness	Kurtosis
30S	2.93E-06	0.00045948	439.3788	193962.296	-7.98032	0.658878	1.499165	9.539247

Table 1.20 Summary of SPY daily realized variance

Table 1.21 suggests that the SPY realized volatilities are extremely right-skewed and leptokurtic. It is very surprising since the daily realized volatilities are calculated by the sum of squared tick-by-tick returns. However, as discussed in the section 1.2.2.2, the tick-by-tick returns are strongly dependent so that the Central Limit Theorem cannot provide accurate approximations in the tick-by-tick data context. In order to explore the relationship between the intraday and daily return dynamics, one must first clearly model the intraday return dynamics. I will further discuss this issue in Chapter 2.

1.6 Conclusions

In this Chapter I explore the unconditional distributions of SPY intraday returns. Two different sampling methods are applied to aggregate the tick-by-tick returns. Although the tick-by-tick returns are largely affected by market microstructural issues, I show that the Central Limit Theorem still can recover the normality to intraday returns. Compared with returns sampled from fixed time interval, returns sampled from fixed number of transactions approaches the normal distribution much faster. For a symbol with ultra-high liquidity like SPY, the intraday returns within each trading day still can be plausibly assumed as normally distributed. This empirical finding is in line with the market microstructural theory of MODM (Harris 1986 and Richardson and Smith 1994).

Besides, the returns sampled from fixed number of tick-by-tick transactions exhibit lower variabilities than the returns sampled from the fixed time interval. This empirical finding implies that the trading intensity should have strong explanatory power over the volatilities of intraday returns. Finally, I use the realized volatility to standardize the SPY daily returns. The standardized SPY daily returns can be well approximated by normal distributions.

These empirical findings might be useful in the field of risk management such as high frequency Value at Risk models. They might also provide some insights regarding the model tick-by-tick price dynamics. From the theoretical perspective, the aggregational normality is closely related the Central Limit Theorem. Nevertheless, due to the large dependence in tick-by-tick returns and effect of market operational details, an explicit explanation based on the parametric modelling of tick-by-tick returns remains as a challenge for the further research.

Chapter 2

Multiplicative Component GARCH Model of Intraday Volatility

2.1 Introduction

The modelling of return dynamics is of both theoretical and empirical importance for financial research. Among the vast literatures on this topic, the ARCH-GARCH models are widely developed and are accepted as the standard techniques for the analysis of return volatility (Engle 1982 and Bollerslev 1986). The ARCH-GARCH models are developed to describe the phenomena of volatility clustering that is the positive autocorrelation in return volatility. However, financial research using high frequency data generally suggest the inadequacy of this standard time series volatility model for the modelling of intraday return dynamics. In general, when applied to high frequency intraday returns, estimations of GARCH models are often contradictory and defy theoretical predictions. For instance, Andersen and Bollerslev (1997) identified that the estimation using GARCH model for returns with different intraday frequencies gave parameters that are inconsistent with the theoretical results on time aggregation of GARCH estimation from Dorst and Nijman (1993).

It is well known that the return volatility systematically varies over the trading day and are highly correlated with the variation of other random variables associated with the transaction records such as the volumes and waiting-time over the trading day. In many markets, this systematical variation of intraday return volatilities is identified as a U-shape pattern in return volatility over the day. Specifically, returns during market opening and closing hours are much more volatile. Such intraday periodic pattern in the return volatility in equity market has been proven to have strong influence on the dynamics of equity returns (Bollerslev 1994, Guillaume et al. 1995). Consequently, in order to carry out meaningful analysis about intraday volatility dynamics, it is necessary to consider the intraday periodic

volatility pattern as a fundamental determinant of the intraday volatility process.

In this chapter, I investigate the difficulties encountered by GARCH volatility models when applied to the high frequency returns. I focus on the strong impact of intraday periodicity of return dynamics on the modelling of intraday volatility process using standard GARCH models. Besides, I construct our estimation of the intraday periodicity in volatility process based on the realized volatility calculated from tick-by-tick transactions. The estimation procedure could be applied to both time-aggregated and transaction-aggregated returns. Based on the discussion of different sampling method of high frequency returns in Chapter 1, the periodicity estimation based on tick-by-tick returns provide the most accurate description of intraday periodic volatility pattern. I also analyze the different impact of the intraday periodic feature of return dynamics on returns sampled from both fixed time interval and fixed number of transactions at various frequencies to verify our estimation of the intraday periodicity of return dynamics. Moreover, to identify the distinctive characteristics of the intraday returns process, the findings are compared with the corresponding features of daily returns series. Finally, I generalize the multiplicative GARCH model for intraday volatility of Engle (2012). I modify the intraday volatility periodicity component of the model to allow it to be conditional on the varying trading period of tick-by-tick transactions.

This chapter is organized as follows. Section 2.2 presents the discussion about literatures regarding the intraday volatility modelling. Section 2.3 describes the data set and presents the intraday return pattern. Section 2.4 exhibits the multiplicative component intraday volatility models. Section 2.5 discusses the specification of the multiplicative volatility model using empirical evidences. It also includes my characterization of intraday volatility periodicity pattern. A relatively simple model that allows the estimation of intraday volatility periodicity to be conditional on the trading period. Section 2.6 contains the empirical analysis regarding our multiplicative component volatility models. I also analyze the statistical properties of the corresponding filtered returns series generated by normalizing the raw return series according to the estimated intraday volatility periodicity and the corresponding estimated daily volatility. I employ high frequency time-aggregated

returns at different frequencies. Estimates of the degree of volatility persistence at the various sampling frequencies are compared with the theoretical aggregation results. Section 2.7 presents the conclusion.

2.2 Literature Review

Early in 1980s, authors like Wood et al. (1985) and Harris (1986) gave the empirical evidence about the distinctive U-shaped pattern in return volatility over the trading day using average intraday returns of stock market. After that, the intraday periodic feature of asset returns is widely reported and considered as a stylized fact regarding financial return series in different markets. For instance, similar evidence for foreign exchange markets can be found in Muller et.al (1990) and Baillie and Bollerslev (1991). Meanwhile, a main topic of financial research using high frequency data is intraday return volatility modelling based on the ARCH-GARCH model of Engle (1982) and Bollerslev (1986). These are partially motivated by an attempt to identify the economic origins of the volatility clustering phenomenon such as the mixture of distributions hypothesis; see for example Clark (1973), Tauchen and Pitts (1983), Harris (1987), Gallant et al. (1991), Ross (1989) and Andersen (1994, 1996).

A direct comparison of these research is difficult because of the different sampling frequencies used by these literatures. Nevertheless, as argued by Ghose and Kroner (1984), Guillaume (1994), Andersen and Bollerslev (1997) and Engle (2012), most of the results regarding degree of volatility persistence and parameter estimation conflicts with the theoretical aggregational results for ARCH-GARCH models by Nelson (1992) and Dorst and Nijman (1993). A plausible explanation is that the strong intraday volatility periodicity cannot be captured by GARCH-ARCH models. Specifically, the intraday volatility periodicity causes the contradictory phenomena between estimation of GARCH-ARCH parameters for returns at different intraday frequencies and theoretical predictions.

In 1990s, efforts on the modelling of intraday volatility periodicity pattern are given by the research group at Olsen and Associates on the foreign exchange market. For instance,

Muller et al. (1990,1993) and Dacorogna et al. (1993) use time invariant polynomial functions to approximate the trading activities in different FOREX exchange market during the 24-hour trading cycle.

Andersen and Bollerslev (1997) gave a more general methodology for the modelling of intraday volatility periodicity pattern. They construct a multiplicative model of daily and intraday volatility for the 5-minute returns on both Deutschemark-dollar exchange rate and US stock market. They express the conditional variance as a product of intraday component and daily component. This specification allows the periodic pattern to be conditional on the current overall level of return volatility. Other closely related models for intraday volatility periodicity can be found in Ghose and Kroner (1996), Andersen and Bollerslev (1998), and Giot (2005). For instance, Giot (2005) adds a deterministic intraday periodicity pattern to the GARCH (1,1) and the EGARCH (Nelson 1991) models and estimates the two model for high frequency return series. Taylor and Xu (1997) provided an alternative specification for high frequency return volatility. They construct an hourly volatility model using an ARCH specification. The conditional variance specification is modified with two elements: the implied volatility and the realized volatility. In 2012, Engle developed a multiplicative component GARCH models based on the work of Andersen and Bollerslev (1998). Specifically, the conditional variance of 10-minute return series in US stock market is expressed as a product of daily, intraday periodicity and stochastic intraday volatility components. Compared with the model of Andersen and Bollerselv (1998) that considers the intraday volatility pattern as deterministic, Engle used two intraday components for the conditional variance model: a deterministic diurnal pattern and stochastic intraday ARCH.

Most of the discussed above use returns sampled from fixed time interval. Consequently, the estimation of intraday volatility periodicity is also based on observations that are equally spaced in time with the same length of interval. Meanwhile, a fundamental feature of tick-by-tick transaction records is that observations are not equally spaced in time. Moreover, as discussed in Chapter 1, returns sampled from fixed number transactions exhibit different statistical characteristics. For instance, the distribution of transaction-aggregated returns approaches the normal distribution much faster than the time-

aggregated returns as the frequencies of returns decrease. Thus, it is interesting to examine whether the models based on time-aggregated returns can be extended to the transaction-aggregated returns, specifically, whether the intraday volatility periodicity estimation based on realized volatility from previous literatures such as Anderson and Bollerslev (1997) can be extended to the transaction-aggregated returns. Finally, I want to estimate the intraday periodicity directly from the tick-by-tick transaction records instead of imposing fixed observing time interval.

Based on the work of Anderson and Bollerslev (1997) and Engle (2012), I use the tick-by-tick transactions to estimate the intraday volatility periodicity component and allow the estimation to be conditional on the trading period during the day. Thus, it can be applied to transaction-aggregated return series and should capture the intraday volatility periodicity more accurately. My estimation method different from the previous research on intraday volatility modelling mainly in three perspectives. First, I generalize the multiplicative component intraday volatility model of Engle (2012) for tick-by-tick return sample. Second, my estimation of intraday periodicity is based on the tick-by-tick transactions, it should give more accurate description on the intraday volatility pattern in the specific time durations. Finally, I test the specification using both time-aggregated and transaction-aggregated returns.

2.3 Data

In this section, I give the data set. In order to comprehensively illustrate the periodic pattern of returns during the trading day, I focus on the time-aggregated return series. In Section 2.3.1, I first exhibit the descriptive statistics about the time-aggregated return series at various frequencies. I focus on the *5-minute* and *10-minute* return series. There are three main reasons for choosing the two frequencies. First, since most of the discussed literatures used these frequencies, I wish to compare the intraday periodic pattern of my data sample with theirs. Second, the preliminary tests show that the intraday pattern is not obvious for returns on higher frequency such as 30s, 1 min and 3 min. Third, the estimation of intraday periodic pattern during short time interval rises computation difficulties. In Section 2.3.2,

I exhibit the intraday periodic pattern of the *5-minute* and *10-minute* return series.

2.3.1 Descriptive statistics

The data set is the tick-by-tick transaction records of symbol SPY during 01/02/2014-12/31/2014. It contains transactions for 252 trading days. The out-of-hour transactions whose occurrence lie outside the time period between 9:30 am and 4:00 pm on each day are removed. Details about the data set and definition of time-aggregated returns can be found in the section 1.2.2.2 in Chapter 1.

	Sample Size	Mean	Variance	Kurtosis	Skewness	Maximum	Minimum
30Sec	196560	1.29618E-07	5.30347E-08	264.8727029	-0.173806698	0.01225414	-0.012983101
1Min	98280	2.59236E-07	9.67974E-08	95.91646403	-0.103161028	0.012458815	-0.012906394
1.5Min	65520	3.88855E-07	1.39691E-07	80.86530506	-0.005621653	0.012036628	-0.013264313
3Min	32760	7.77709E-07	2.70802E-07	39.85096976	-0.238157281	0.011012149	-0.014459855
5Min	19656	1.29618E-06	4.06658E-07	8.436150587	0.075305595	0.006182118	-0.005359357
10Min	9828	2.59236E-06	8.03275E-07	7.918815602	-0.037720742	0.008269006	-0.006320998
15Min	6552	3.88855E-06	1.17727E-06	7.388130566	-0.079263515	0.007751481	-0.006653523
30Min	3276	7.77709E-06	2.34033E-06	6.660810707	-0.095190193	0.011082252	-0.009269532
78Min	1260	2.02204E-05	6.20987E-06	6.687141886	-0.195275515	0.014103076	-0.014359004

Table 2.1 Descriptive statistics of time-aggregated returns

A detailed discussion of the descriptive statistics of time-aggregated return series at various frequencies is provided in Section 1.2.2.2 in Chapter 1. In this section I focus on the 5-minute and 10-minute return series. It seems natural to choose the time-aggregated return to present an illustration of intraday periodic pattern of the returns over the trading day. I choose the absolute value of the returns and squared returns as the proxies of the return volatility.

	Mean	Variance	Kurtosis	Skewness	Maximum	Minimum	First Lag autocorrelation
5-Minute Return	1E-06	4.07E-07	8.4362	0.0753	0.006182	-0.0054	-0.015058796
Absolute 5-Minute Return	0.0004	2.13E-07	14.733	2.5577	0.006182	0	0.256396005
Squared 5-Minute Return	4E-07	1.23E-12	237.48	11.289	3.82E-05	0	0.179436501

Table 2.2 Descriptive statistics of 5-minute return series, absolute 5-minute return series and squared 5-minute return series

	Mean	Variance	Kurtosis	Skewness	Maximum	Minimum	First Lag autocorrelation
10-Minute Return	2.59E-06	8.03E-07	7.918815602	-0.03772074	0.008269006	-0.006321	-0.02467549
Absolute 10-Minute Return	0.000621954	4.16E-07	13.50275706	2.470604056	0.008269006	0	0.247384273
Squared 10-Minute Return	8.03E-07	4.46E-12	197.0880086	10.11289837	6.84E-05	0	0.16996921

Table 2.3 Descriptive statistics of 10-minute return series, absolute 10-minute return series and squared 10-minute return series

The mean of the 5-minute return series and 10-minute return series are 1E-06 and 2.59E-06 respectively. Given their variance of 4.07E-07 and 8.03E-07, the mean of 5-minute return series and 10-minute return series can safely assumed to be zero. However, as suggested by the kurtosis in Tables 2.2 and 2.3, they are clearly not normally distributed. An interesting fact is that although the 5-minute return series have a positive skewness of 0.0753, the 10-minute return series have a negative skewness of -0.0378. Meanwhile, the maximum and minimum of the 5-minute returns are 0.0062 and -0.0054 respectively. These values do not suggest of sharp discontinuities in the return series. The maximum and minimum of 10-minute return series exhibit similar characteristics.

Note that the first lag autocorrelations of the 5-minute and 10-minute returns, -0.0151 and -0.0247 respectively, are significant since corresponding lower confidence bounds at the significance level of 5% are -0.0143 and -0.0201 respectively. The values of the correlation are nearly negligible from the economic perspective. In contrast, the volatilities of the 5-minute returns and 10-minute returns have highly significant positive autocorrelations. For instance, the absolute 5-minute return series has a first lag autocorrelation of 0.2564 and the absolute 10-minute return series has a first lag autocorrelation of 0.2474. The upper confidence bounds of the 5-minute return series and the 10-minute return series at the significance level of 5% are 0.0143 and 0.0201 respectively. The positive correlations of return volatility clearly suggest the phenomena of volatility clustering. Detailed discussion of time-aggregated returns autocorrelations at different frequencies can be found in Chapter 1. The squared returns exhibit slightly smaller but also highly significant positive first lag autocorrelations.

2.3.2 Intraday return periodicity

In order to evaluate the intraday volatility periodicity, I present the plot of average and absolute average during the trading day for both 5-minute return series and the 10-minute return series first.




Figure 2.1 The average SPY 5-minute return over the trading day

The average SPY 5-minute returns are centered around 0 but fluctuate dramatically during the trading day. Figure 2.1 suggests that the average 5-minute returns during the opening and closing hours of the trading day are much more volatile. For instance, three of the four violations of the constant 5% confidence band for the null hypothesis of an *iid* series occur during the first and last ten intervals. Specifically, the first ten interval represents the period from 9:30 to 10:20 and last ten interval represents the period from 3:10 to 4:00. The sharp drop off during such intervals provide graphic evidence. Also note that, the Andersen-Darling statistics cannot reject the null hypothesis that the average 5-minute returns come from a normal distribution at significance level of 5%.




Figure 2.2 The average SPY 10-minute return over a trading day

Figure 2.2 gives similar but more clear presentation of intraday return periodicity for the average 10-minute return series. Three of the four violations of the constant 5% confidence band for the null hypothesis of an *iid* series occur during the first and last five intervals, which is consistent with the finding of the average 5-minute return series. As the case of average 5-minute returns, the average 10-minute returns also can be safely assumed to be normally distributed. It is worthy to mention that during the end of trading day, the dramatic variation of SPY 5-minute returns in Figure 2.1 no longer exist in Figure 2.2. In contrast, the 10-minute SPY average returns tend to be more volatile than 5-minute SPY average return during the opening hours.

Figure 2.3 exhibits the volatility (measured as the absolute average returns) of SPY average 5-minute returns during the trading day. It clearly verifies the pronounced U-shape of return volatility in financial markets. The volatility ranges from a low of around 0.03% at the 13:05 to a high of around 0.08% at the 9:35. It starts out at a relatively high level of 0.07% at 9:30 and decays slowly to 0.03% at 13:05, and then it again surges to around 0.05% at 16:00. The highly significant first lag positive autocorrelation of 0.256 in Table 2.2 also support the existence of intraday periodic volatility pattern. Figure 2.4 suggests similar periodic pattern for 10-minute return volatility. It also can be concluded from Figures 2.3 and 2.4 that the returns are more volatile during the opening hours than the closing hours, which corresponds to the 10am FED news announcements.




Figure 2.3 The absolute average SPY 5-minute return over a trading day

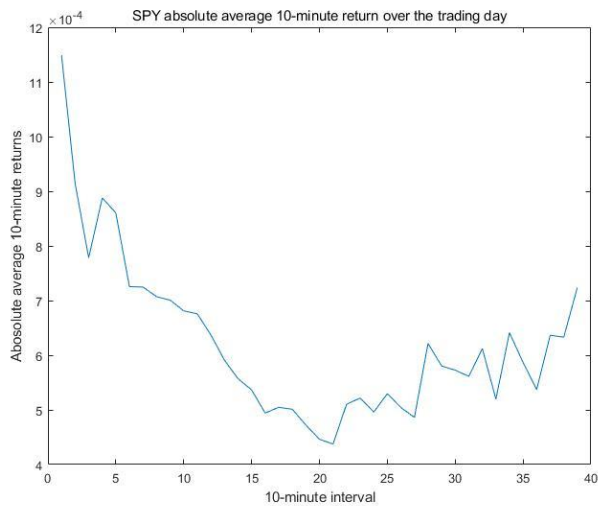


Figure 2.4 The absolute average SPY 10-minute return over a trading day

To support my prior investigation of the unconditional intraday volatility periodic pattern, I now examine the autocorrelations of the absolute 5-minute and absolute 10-minute return series to explore the dynamic feature of those two return series.

	5-minute return	absolute 5-minute return	10-minute return	absolute 10-minute return
Lag1	-0.015058796	0.256396005	-0.024675493	0.247384273
Lag2	-0.02727115	0.27393497	-0.007165027	0.262341418
Lag3	0.011513466	0.266731377	0.025501631	0.271487025
Lag4	-0.002240064	0.25614196	-0.018678124	0.256228375
Lag5	-0.001991967	0.248556941	0.038864765	0.246937983
Lag6	0.007446405	0.267334076	0.029325701	0.236207573

Lag7	0.005399493	0.265061453	0.00940831	0.230465298
Lag8	-0.004431232	0.230266141	-0.002262395	0.217919076
Lag9	0.004174203	0.244515482	-0.002524387	0.209070546
Lag10	0.023174997	0.230559214	0.02308969	0.199440976
Lag11	0.022193775	0.232530198	0.008302441	0.189182483
Lag12	0.00899738	0.219332721	-0.015269527	0.188770108
Lag13	-0.005247707	0.217748927	-0.011524974	0.185738101
Lag14	0.013262886	0.222211507	0.025102784	0.194993096
Lag15	0.007466957	0.219513545	0.009805094	0.174629519
Lag16	-0.009635652	0.199277327	0.004593636	0.175036007
Lag17	-0.004045313	0.214672245	-0.027126605	0.155205843
Lag18	0.006474428	0.200850752	-0.029243554	0.152888546
Lag19	0.00035781	0.184185116	-0.017268827	0.145260443
Lag20	0.009019369	0.18152872	0.010949894	0.141655001
Upper Bounds	0.01426535	0.01426535	0.020174251	0.020174251
Lower Bounds	-0.01426535	-0.01426535	-0.020174251	-0.020174251

Table 2.4 Autocorrelations of average returns and absolute average returns

As discussed previously, the 5-minute return series and 10-minute return series can be plausibly considered as uncorrelated. For instance, although 5-minute return has significant first and second lag autocorrelations, the values of the autocorrelations are as small as -0.015 and -0.027 respectively. In contrast, the absolute return series exhibits highly significant strong positive correlations at almost all 20 lags. Observe that the autocorrelation of both absolute 5-minute and absolute 10-minute return series decay very slowly. If there is a periodic pattern of the volatility during each trading day, the correlations should also have a 24-hour cycle pattern.

To further investigate the intraday volatility periodicity pattern, I plot the correlogram for both absolute 5-minute return and absolute 10-minute return series for up to five days.




Figure 2.5 Five days SPY absolute 5-minute return correlogram

The Figures 2.5 and 2.6 present autocorrelation pattern of the absolute returns. The intraday volatility periodicity is exhibited as the corresponding U-shape in the correlogram. The correlations at daily frequency decay very slowly during the five days. Observe that the slowly declining U-shape cycles around every 80 intervals for the 5-minute lag and every 40 intervals for the 10-minute lag. Since the NYSE operates from 9:30 to 4:00 on each trading day, U-shape cycles on daily frequency exactly.




Figure 2.6 Five days SPY absolute 10-minute return correlogram

2.4 Multiplicative component intraday volatility model

In this section I present my multiplicative component intraday volatility model. The model is based on the work of Andersen and Bollerslev (1997) and Engle (2012). In order to illustrate the reasoning behind the model specification, I first discuss the interpretation of the intraday volatility periodicity pattern in section 2.4.1. I then present the model and discuss the econometric issues in section 2.4.2.

2.4.1 Interpretation of the intraday return volatility

The periodic pattern of the intraday volatility discussed in Section 2.3 suggests that a direct use of the ARCH-GARCH for modelling of the intraday return volatility could be inappropriate. The main reason is that the standard ARCH-GARCH models impose a geometric decay for the return autocorrelation structure. It cannot capture the periodic pattern of the volatility correlation that are reported in Section 2.3. In order to jointly describe the intraday periodicity and daily conditional heteroskedasticity, Andersen and Bollerslev (1997) present the following stylized specification of the intraday returns. Consider an intraday time-aggregated logarithm return series $\{r_{t,n}\}_{t,n}$, where $r_{t,n}$ represent the n th return at day t . Suppose there are N observation during each trading day. Then

$$r_{t,n} = \sigma_t \frac{1}{\sqrt{N}} s_n z_{t,n}, \quad (2.1)$$

where σ_t is the conditional daily volatility at day t , s_n is the intraday volatility periodicity during the period of $r_{t,n}$ and $z_{t,n}$ is an *iid* series of zero mean and unit variance that is independent of the daily volatility σ_t and the intraday periodicity s_n . The volatility components σ_t and s_n are required to be positive. For instance, $\sigma_t > 0$ for all t and $s_n > 0$ for all n . The intraday volatility periodicity component s_n is a periodic function at daily frequency. Specifically, $s_{n+jN} = s_n$ for all j and n . Thus, the daily return at day t R_t , is

$$R_t = \sum_{n=1}^N r_{t,n} = \sigma_t \frac{1}{\sqrt{N}} \sum_{n=1}^N s_n z_{t,n}. \quad (2.2)$$

With respect to the model estimation, Andersen and Bollerslev (1997,1998) apply a flexible Fourier transform to estimate the intraday volatility periodicity component s_n and use standard GARCH volatility model to estimate the conditional daily volatility σ_t . For

further details of the Fourier transform, see the appendix of Andersen and Bollerslev (1997).

Under the assumption that the intraday returns are serially uncorrelated, the daily conditional variance is the sum of variances from each time interval. Thus,

$$E \left(\sum_{n=1}^N \frac{r_{t,n}^2}{\sigma_t^2} \right) = 1. \quad (2.3)$$

(2.3) immediately gives that $\sum_{n=1}^N s_n^2/N = 1$. This intuitive specification of intraday volatility has two important features. First, it extends naturally to standard specifications of daily return volatility models with intraday innovations and deterministic daily volatility. Specifically, in the case that $s_n = 1$ for all n , the $r_{t,n} = \sigma_t \frac{1}{\sqrt{N}} z_{t,n}$. Consequently, $r_t = \sigma_t \frac{1}{\sqrt{N}} \sum_{n=1}^N z_{t,n}$. Thus, the daily return at day t is the sum of product of independent intraday innovations and daily conditional volatility patterns. The innovation component $\frac{1}{\sqrt{N}} \sum_{n=1}^N z_{t,n}$ is clearly an *iid* series with zero mean and unit variance.

Besides, the model does not affect the autocorrelation structure of the returns measured at daily frequency. When returns are measured at daily frequency, R_t the daily return at day t is $R_t = \sum_{n=1}^N \frac{s_n}{\sqrt{N}} z_t \sigma_t$ where z_t is an *iid* series of zero and unit variance. Thus, the expected absolute return $E|R_t|$ is,

$$E|R_t| = \sum_{n=1}^N \frac{s_n}{\sqrt{N}} E|z_t| \sigma_t. \quad (2.4)$$

Observe that since $E(\sum_{n=1}^N s_n^2/N) = 1$, $\sum_{n=1}^N \frac{s_n}{\sqrt{N}} \geq 1$ since $\left(\sum_{n=1}^N \frac{s_n}{\sqrt{N}} \right)^2 = (\sum_{n=1}^N s_n)^2/N \geq \sum_{n=1}^N s_n^2/N$. Consequently, the expected daily absolute return is an increasing function of the variation of the intraday periodicity pattern. Let $c = (E|z_t|)^{-2} - 1 > 0$. I immediately have,

$$Corr(|R_t|, |R_\tau|) = \frac{Cov(\sigma_t, \sigma_\tau)}{Var(\sigma_t) + cE(\sigma_t^2)}. \quad (2.5)$$

(2.5) clearly suggests that the model does not impose any structural changes to the autocorrelation of the daily returns.

Second, the model gives a qualitative description of the impact of intraday volatility

periodicity on the autocorrelation of the absolute intraday returns. Specifically, suppose that $n \geq m$, then

$$Corr(|R_{t,n}|, |R_{\tau,m}|) = \frac{\sum_{j=1}^N s_j s_{j-(n-m)} Cov(\sigma_t, \sigma_\tau) + Cov((s_n, s_m) E^2(\sigma_t))}{\frac{Var(\sigma_t) (\sum_{j=1}^N s_j)^2}{N} + c E(\sigma_t^2) \frac{(\sum_{j=1}^N s_j)^2}{N} + E^2(\sigma_t) Var(s)}. \quad (2.6)$$

Notice that, when $n = m$, equation (2.6) reduces to

$$Corr(|R_{t,n}|, |R_{\tau,n}|) = \frac{Cov(\sigma_t, \sigma_\tau) + (\frac{Var(s)}{(\sum_{j=1}^N s_j)^2}) E^2(\sigma_t)}{Var(\sigma_t) + c E(\sigma_t^2) + (\frac{Var(s)}{(\sum_{j=1}^N s_j)^2}) E^2(\sigma_t)}. \quad (2.7)$$

(2.6) gives a qualitative description between the intraday periodicity and the daily heteroskedasticity of the absolute returns. The positive covariance of the daily return volatility $Cov(\sigma_t, \sigma_\tau)$ is strong when t and τ are very close. Consequently, it causes positive dependence in the absolute returns, as the distance between t and τ grows larger this effect becomes less significant. This implication is consistent with the slow decay of the correlations in Figures 2.5 and 2.6.

At the same time, the strong intraday volatility periodicity also has an impact on the correlations of absolute returns. For instance, the autocorrelations of 5-minute return series reaches the lowest value generally at the forty lags, which is exactly half of the trading day. The covariance of the intraday volatility periodicity $Cov((s_n, s_m))$ is expected to reach the minimal at the same time according to the U-shape. Figure 2.6 suggests that the 10-minute return series exhibit similar characteristics as well. These findings further confirm the correspondence between the qualitative implications of (2.1) and the autocorrelation structure of the intraday absolute returns.

Although the model of (2.1) gives a plausible explanation of the autocorrelation structure of the absolute intraday returns and might serve as a starting point for the high frequency volatility modelling, the simplistic intraday volatility specification $\sigma_t \frac{1}{\sqrt{N}} s_n$ implies that the only intraday component in the intraday volatility pattern is s_n .

Based on the intraday return volatility model of Andersen and Bollerslev (1997,1998), Engle (2012) proposed the following multiplicative specification for intraday return volatility to provide a more realistic volatility dynamic description,

$$R_{t,n} = \sigma_t \epsilon_{t,n} s_n z_{t,n}, \quad (2.8)$$

where σ_t is the conditional volatility at day t , s_n is the intraday volatility periodicity component, $\epsilon_{t,n}$ is the stochastic intraday volatility component with $E(\epsilon_{t,n}^2) = 1$ and $z_{t,n}$ is an *iid* series that follows the normal distribution with zero mean and unit variance. Moreover, the normalized return $y_{t,n} = \frac{r_{t,n}}{\sigma_t s_n}$ is assumed to follow a GARCH (1,1) process. Specifically,

$$\begin{aligned} y_{t,n} | F_{t,n-1} &\sim N(0, \epsilon_{t,n}), \\ y_{t,n} &= \frac{r_{t,n}}{\sigma_t s_n} = \epsilon_{t,n} z_{t,n}, \\ \epsilon_{t,n}^2 &= \omega + \alpha y_{t,n-1}^2 + \beta \epsilon_{t,n-1}^2. \end{aligned}$$

Similarly, given the assumption that the intraday returns are not serially correlated, one immediately has $E\left(\sum_{n=1}^N \frac{r_{t,n}^2}{\sigma_t^2}\right) = 1$. Then $E\left(\sum_{n=1}^N \frac{r_{t,n}^2}{\sigma_t^2}\right) = \sum_{n=1}^N s_n^2 = 1$. Regarding the model estimation, for the intraday volatility periodicity pattern s_n , Engle applies a commercially available risk measure to estimate σ_t and further the calculates the s_n as the variance of returns in each time interval after deflating by the conditional daily variance σ_t .

Specifically,

$$E\left(\frac{R_{t,n}^2}{\sigma_t^2}\right) = E(\epsilon_{t,n}^2 z_{t,n}^2 s_n) = s_n. \quad (2.9)$$

Note that the model of both (2.1) and (2.8) are based on the return series that are sampled from fixed time interval. In order to study the intraday volatility periodicity for the transaction-aggregated return series, I must estimate the intraday volatility periodicity for irregular spaced time interval during the trading day.

2.4.2 Multiplicative component intraday volatility

In this section I present the multiplicative component intraday volatility model. My model is based on the Andersen and Bollerslev (1997, 1998) and Engle (2012). In 1997, Andersen and Bollerslev proposed the intraday volatility model of (2.1). They expressed the daily

conditional variance as the product of intraday and daily volatility. As discussed in section 2.4.1, this specification successfully describes the periodic structure of the absolute intraday return and can be extended to standard daily volatility model when returns are measured at daily frequencies. This prior specification is proved to be very fruitful. In 1998, Andersen and Bollerslev further add a component to the multiplicative specification of intraday volatility to capture macroeconomic announcements. The model is widely accepted for the modelling of intraday volatility in foreign exchange data. Later in 2012, Engle propose a multiplicative intraday volatility model that express the intraday conditional variance as a product of intraday periodicity, intraday variance and daily variance.

When using tick-by-tick data for an equity with high liquidity like SPY, the component that accounts for macroeconomic announcements in the volatility model is not very practical. E.g., Andersen and Bollerslev (1998) and Rangel (2009). First, most important macroeconomic announcements such as the release of new monetary policy happen before the stock market opens. Second, the timing of announcements that are particularly important to equity markets are hard to predict. Third, it is very difficult to capture and measure the market response to the macroeconomic announcement in a circumstance that the duration between two consecutive transactions is less than 0.1s. Finally, asymmetric information and microstructure are generally believed to play a decisive role in the high frequency return dynamics. However, macroeconomic or public announcement dummies cannot account for information arrival through order flow (Engle 2012).

Besides, the intraday volatility models of (2.1) and (2.8) are based on returns sampled from fixed time interval. It cannot be applied directly for the transaction-aggregated return series. For instance, consider a tick-by-tick return series $r_{t,n}$, where $r_{t,n}$ represents the n th tick-by-tick return at day t . In general, $w_{t,n} \neq w_{\tau,n}$ when $t \neq \tau$ where $w_{t,n}$ is the waiting-time between transaction n and transaction $n+1$ at day t . Consequently, the periodic component s_n in (2.1) and (2.8) should be conditional on both t and n . I therefore generalize the model of (2.8) by Engle (2012) to allow the intraday periodicity component to be conditional on the corresponding waiting-time of $r_{t,n}$. I present our intraday volatility model as the follows.

$$r_{t,n}/\sqrt{w_{t,n}} = \sigma_t \epsilon_{t,n} s_{t,n} z_{t,n}, \quad (2.10)$$

where σ_t is the conditional volatility at day t , $w_{t,n}$ is the duration that corresponds to $r_{t,n}$, $s_{t,n}$ is the intraday volatility periodicity component, $\epsilon_{t,n}$ is the stochastic intraday volatility component with $E(\epsilon_{t,n}^2) = 1$ and $z_{t,n}$ is an *iid* series that follows the normal distribution with zero mean and unit variance. $\epsilon_{t,n}$ and $z_{t,n}$ are assumed to be independent of each other. σ_t , $s_{t,n}$ and $\epsilon_{t,n}$ are strictly positive. That is, $\sigma_t > 0$ for all t , $s_{t,n} > 0$ for all t, n , $\epsilon_{t,n} > 0$ for all t, n .

It is noticeable that in the specification of (2.10), I do not consider the duration $w_{t,n}$ as a random variable. It is given exogenously. The incorporation of $w_{t,n}$ in (2.10) therefore is for the normalization of irregularly spaced returns.

There are several reasons behind this consideration. First, joint modelling of durations and volatilities requires an explicit parametric modelling of the arrivals of transactions that is beyond the scope of this chapter. I will elaborate the duration modelling in the ACD framework in Chapter 3. Second, the intraday periodic pattern of volatilities and durations are quite different. Adding the intraday periodic component into (2.10) would further complicate the model specification. Finally, the durations are in general correlated with volatilities. The modelling of durations often involves explanatory variables that are volatility measures. Hence, a multiplicative specification seems to be over simplistic.

Under (2.10) the daily return R_t at day t is,

$$R_t = \sigma_t \sum_{n=1}^{N_t} \sqrt{w_{t,n}} \epsilon_{t,n} s_{t,n} z_{t,n}, \quad (2.11)$$

where the N_t represents the total number of tick-by-tick transaction records at day t . It is clear that $E(R_t) = 0$. Moreover, the expected value of the squared daily return $E(R_t^2)$ is,

$$E(R_t^2) = \sigma_t^2 \sum_{n=1}^{N_t} E(w_{t,n} \epsilon_{t,n}^2 s_{t,n}^2 z_{t,n}^2) = \sigma_t^2 \sum_{n=1}^{N_t} w_{t,n} s_{t,n}^2. \quad (2.12)$$

Consequently, $\sum_{n=1}^{N_t} w_{t,n} s_{t,n}^2 = 1$. Without the presence of the intraday periodicity, in which case $s_{t,n} = \frac{1}{\sqrt{N_t}}$ for all n , the daily return R_t is

$$R_t = \sigma_t \sum_{n=1}^{N_t} \frac{1}{\sqrt{N_t}} \sqrt{w_{t,n}} \epsilon_{t,n} z_{t,n}. \quad (2.13)$$

Observe that $\sum_{n=1}^{N_t} \frac{1}{\sqrt{N_t}} \sqrt{w_{t,n}} \epsilon_{t,n} z_{t,n}$ is an independent but not necessarily identical distributed sequence of zero mean and unit variance. I further model the $\epsilon_{t,n}$ by a GARCH (1,1) process, that is

$$\begin{aligned} y_{t,n} | F_{t,n-1} &\sim N(0, \epsilon_{t,n}), \\ y_{t,n} &= \frac{r_{t,n}}{\sigma_t s_{t,n} \sqrt{w_{t,n}}} = \epsilon_{t,n} z_{t,n}, \\ \epsilon_{t,n}^2 &= \omega + \alpha y_{t,n-1}^2 + \beta \epsilon_{t,n-1}^2 \end{aligned} \quad (2.14)$$

In a summary of the extensive ARCH/GARCH literatures, Bollerslev, Chou, and Kroner (1992) argued that GARCH (1,1) is the most popular model. Moreover, even if the conditional variance specification requires more lag terms, the models are not of a higher order than GARCH (1,2) or GARCH (2,1). In 1990, Nelson gave an explanation of this empirical result. Specifically, efficiency considerations favor models of a lower order since many ARCH/GARCH models could be consistent filters of a particular diffusion process.

With respect to the model estimation, Engle (2012) used a commercially available daily volatility measure that is based on the multifactor risk model of Fabozzi, Jones, and Vardharaj (2002). The use of a daily variance forecasts from the structural analysis allows the intraday volatility model to capture industry factor and common liquidity factors. Nevertheless, the market efficiency theory asserts that such information should be included in the market price. Moreover, given the strong ARCH effect of our daily return of SPY, a GARCH modelling seems to be natural. Thus, I capture the σ_t by a GARCH specification. Note that the daily variance component σ_t can be estimated by daily realized volatility approaches such as the Engle and Gallo (2006) and Zhang et.al. (2005).

For the intraday periodicity component $s_{t,n}$, Andersen and Bollerslev (1997,1998) applied a flexible Fourier transform to estimate the intraday volatility periodicity component $s_{t,n}$. Their approach has an important advantage that it allows the $s_{t,n}$ to be conditional on t . Specifically, the intraday periodic pattern is time varying and is dependent on the conditional daily volatility σ_t . However, although the periodic pattern of financial returns is a well-known empirical stylized fact, there is no widely accepted economic theory that

could stipulate the parametric specification of dynamic structure of the intraday periodicity. As a result, nonparametric procedure seems natural. Thus, I follow the methodology of Engle (2012) and consider $s_{t,n}$ as the variance of returns during its corresponding waiting-time deflated by the daily conditional variance. This specification allows the daily conditional variance to be completely free. Specifically, observe that the specification of (2.10) implies that

$$E\left(\frac{r_{t,n}^2}{\sigma_t^2}\right) = E(w_{t,n}\epsilon_{t,n}^2 s_{t,n}^2 z_{t,n}^2) = w_{t,n} s_{t,n}^2. \quad (2.15)$$

(2.15) establishes the basis of the estimation of intraday periodicity pattern.

2.5 Intraday volatility components and model estimation

In this section I further explain the reasoning behind of the specification of (2.10) and discuss the econometrics regarding the estimation of the model. In Section 2.5.1 I discuss the ARCH effect of the daily returns of SPY and the impact of the ARCH effect on the intraday return series. I also present my modelling for the daily conditional variance σ_t . In Section 2.5.2 I present my estimation methodology for the intraday volatility periodicity. In Section 2.5.3 I discuss the econometrics regarding my estimation.

2.5.1 ARCH effect of daily returns

In this section I discuss the ARCH effect of daily returns and estimate the daily conditional volatility σ_t . I first present the descriptive statistics of the SPY daily returns from January 3, 2005 to December 31, 2014. There are total 2517 daily observations and overnight effect are not included.

	Mean	Variance	Kurtosis	Skewness	Maximum	Minimum	First Lag autocorrelation
Daily returns	2.03E-06	0.000104	14.64993	-0.52421	0.076669	-0.09421	-0.08628
Absolute daily returns	0.006581	6.02E-05	25.97425	3.688621	0.094207	0	0.330214
Squared daily returns	0.000103	1.46E-07	197.7099	11.9846	0.008875	0	0.227522

Table 2.5 Descriptive statistics of daily return series, absolute daily return series and squared daily return series




Figure 2.7 SPY daily returns from January 3, 2005 to December 31, 2014

Figure 2.7 suggests that the daily returns between October 2008 and December 2008 are extremely volatile, which corresponds to the burst of the subprime mortgage financial crisis in 2008. The other busy period following since the end of 2009 corresponds to the European debt crisis. In general, the daily returns fluctuate dramatically around zero. This behavior gives a graphic illustration of the conditional heteroscedasticity.

Besides, although the mean of the daily SPY returns can be safely assumed to be zero, the kurtosis of 14.65 and skewness of -0.52 suggest that the returns are not normally distributed. The daily returns have a small but significant negative first lag autocorrelation. Given the standard significance level of 5%, the lower bound and the upper bound are -0.03986 and 0.03986 respectively. As expected, the daily volatility presents strong volatility clustering phenomenon. The first lag autocorrelations of the absolute daily and squared daily returns are of nonnegligible values and are highly significant. The following correlograms of the absolute daily returns and the squared daily returns provide further evidence.

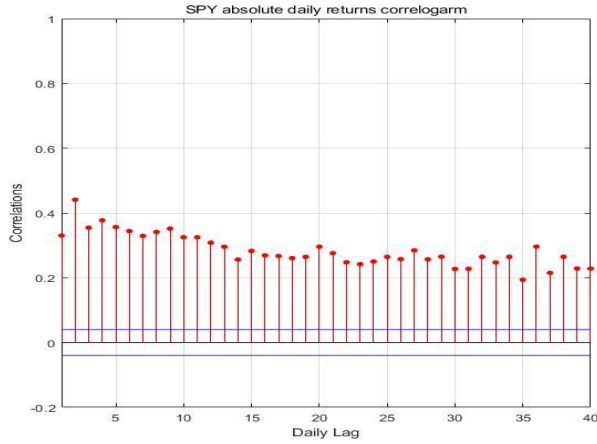


Figure 2.8 SPY absolute daily return correlogram

Figure 2.8 exhibits the slow decay of the correlation in the volatility of SPY daily returns.




Figure 2.9 SPY squared daily return correlogram

To test the ARCH effect statistically, I apply the ARCH test developed by Engle (1988). Consider an ARCH (p) process,

$$y_t^2 = a_0 + a_i \sum_{i=1}^p y_{t-i}^2 + \varepsilon_t.$$

The alternative hypothesis is that there is at least one $a_i \neq 0$ for $i=1,2,3,\dots, q$. The test statistic is the Lagrange multiplier statistic TR^2 where T is the sample size and R^2 is obtained from fitting the ARCH(p) model via regression. Under the null hypothesis that $a_i = 0$ for all i , the asymptotic distribution of the test statistic TR^2 is Chi-Square with

p degrees of freedom.

A commonly accepted procedure for choosing the number of lags in the ARCH modelling is that one first estimates different lag structures and chooses the one with the minimal Akaike Information Criterion statistics (AIC). However, this procedure in practice usually favors a relatively long lag structure. Thus, I test the null hypothesis that there is no ARCH effect in the daily series against the alternative hypothesis of an ARCH model with two lagged squared innovations that is equivalent to a GARCH (1,1) model locally (Bollerslev 1986). Although the GARCH (1,1) model is not necessarily the preferred model, it still gives a simple and comprehensive approximation to the dependence structure in the autocorrelation of squared returns. The result shows that the test rejects the null hypothesis at the significance level of 1% with a p value of zero for the test statistics. Specifically, the value of the test statistics is 747.5967 and the critical value for the given significance level of 1% is 9.2103.

The pronounced ARCH effect in the previous discussion might not be very surprising since it is widely documented for financial research. However, the question remains that how this ARCH effect observed from daily returns, the aggregation of tick-by-tick returns, influences the intraday volatility process. According to my prior investigation of the ARCH effect on the daily return series, an intraday volatility process with conditional heteroskedasticity seems natural. Nevertheless, the poor performance of ARCH-GARCH models for the modelling of intraday volatility also is widely reported. (Cumby et al. 1993, West and Cho 1995, Figlewski 1995 and Jorin 1995). In order to explicitly explain the intraday volatility process, I first investigate the relation between the daily observed ARCH effect and intraday volatility pattern.

To assess the influence of the ARCH effect observed from daily returns on the intraday returns, we explore the relation between one-step-ahead volatility estimation provided by GARCH models on daily returns and other daily ex post return variability measures calculated from intraday returns. I choose the daily returns from January 3, 2005 to December 31, 2013 to estimate the GARCH model and further compare the one-step-ahead

GARCH estimations with volatility measures calculated from intraday returns using sample from January 2, 2014 to December 31, 2014.

For the conditional heteroskedasticity modelling of the daily returns, I choose the MA (1)-GARCH (1,1) model specification. The moving average term MA (1) is included to explain the weak but significant negative first order autocorrelation in Table 2.5. The model is specified as,

$$\begin{aligned} R_t &= u + \theta \varepsilon_{t-1} + \varepsilon_t, \\ \varepsilon_t &= \sigma_t z_t, \\ \sigma_t^2 &= \omega + \alpha \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2, \end{aligned} \quad (2.16)$$

where z_t is an *iid* sequence and $z_t \sim N(0,1)$. The ARMA is proposed by the thesis of Peter Whittle (1951). The maximum likelihood estimation regarding the ARMA-GARCH process can be found in Ling and Li (1997,1998), McAleer (2003) and Franco (2004). Note that, the GJR-GARCH could be another natural candidate for modelling σ_t^2 since the daily return series has a skewness of -0.5241. The leverage effect that is the asymmetric contribution of returns with different signs to the variability of the prices is expected to be pronounced when returns are measured at daily frequency.

	Values	Standard Error	T statistics	P values
u	0.000272	0.000137	1.988748	0.046729
θ	-0.06674	0.024112	-2.76797	0.005641
ω	1.69E-06	5.11E-07	3.311307	0.000929
β	0.873353	0.010796	80.89544	0
α	0.104926	0.009394	11.16934	5.76E-29

Table 2.6 Parameter estimation of the MA (1)-GARCH (1,1) model

Table 2.6 presents the estimated MA (1)-GARCH (1,1) model. Observe that α is relatively large and β is relatively small. The sum of β and α is around 0.978 that represents a high level of volatility persistence. The following Figure 2.10 exhibits the plot of standardized residuals and conditional daily variances. Observe the sudden surge of the conditional variance between the interval between 900 to 1000 that corresponds to the dramatic variation of daily returns during the same period in Figure 2.7.




Figure 2.10 The conditional variance and standardized residuals of MA (1)-GARCH (1,1) fitting




Figure 2.11 The correlogram of squared standardized residuals of MA (1)-GARCH (1,1) fitting

Figure 2.11 exhibits the autocorrelation function of the squared standardized residuals. The autocorrelations are eliminated as expected. Although the MA (1)-GARCH (1,1) captures the conditional heteroskedasticity successfully as suggested by Figures 2.10, 2.11, Figure 2.12 suggests the standardized residuals still presents fat tails. Another interesting finding is that, compared with Figure 2.7, the volatility of SPY daily returns in Figure 2.10 is much regular. Specifically, the effect of financial crisis in 2008 still stays but other noisy periods start to drop. It might suggest the different effect of intraday events such as mini flash crash and events lasting over days such as the financial crisis in 2008 over the daily volatility pattern of SPY.




Figure 2.12 The Quantile-Quantile Plot of the standardized residuals

I now turn to the ex post return variability measure that are calculated from intraday returns. The two measures that we choose are the realized variance and the cumulative absolute returns. Specifically, $\sum_{n=1}^{N_t} |r_{t,n}|$ and $\sum_{n=1}^{N_t} r_{t,n}^2$.

As discussed in Chapter 1, compared with volatility measures that are based on daily returns that cannot capture the variability of intraday prices movement, the realized volatility gives a more realistic estimation since it takes, at the highest frequency, each tick-by-tick observation into account. For instance, the intraday price can fluctuate dramatically over the trading day but ends with a value that is very close to the opening price.

We present the comparison of the one-step-ahead MA (1)-GARCH (1,1) daily forecasts with the absolute daily return $|R_t| = [\sum_{n=1}^{N_t} r_{t,n}]$, realized variance $\sum_{n=1}^{N_t} r_{t,n}^2$ and cumulative absolute returns $\sum_{n=1}^{N_t} |r_{t,n}|$. Daily returns of SPY from January 3, 2005 to December 31, 2013 are employed to estimate the MA (1)-GARCH (1,1) model. Tick-by-tick transactions from January 2, 2014 to December 31, 2014 are used to calculate the realized volatility and cumulative absolute returns and are aggregated to daily frequency to serve as the out-sample for the GARCH forecasts. All series are normalized to have an average of one.

The following Figure 2.13 presents the one-step-ahead GARCH (1,1) forecasts and the daily absolute returns $|R_t|$. The daily absolute returns fluctuate arbitrarily around the GARCH predictions. It seems intuitively that the GARCH forecasts should have weak explanatory power of the observed daily variability of returns that is measured by, $|R_t| = [\sum_{n=1}^{N_t} r_{t,n}]$. Further, it might also implicitly suggest that the ARCH effects of the daily returns have negligible impact on the intraday volatility. However, the suggestion is misleading, since

$$|R_t| = [\sum_{n=1}^{N_t} r_{t,n}] = \left| \sum_{n=1}^{N_t} \ln \left(\frac{P_{t,n+1}}{P_{t,n}} \right) \right| = |\ln P_{t,N_t+1} - \ln P_{t,1}|. \quad (2.17)$$

(2.17) suggests that under the assumption that there is no over night effect $|R_t| = [\sum_{n=1}^{N_t} r_{t,n}]$ should coincide with the daily absolute returns. In real world, the two concepts are not equivalent for sure. Thus, $|R_t| = [\sum_{n=1}^{N_t} r_{t,n}]$ only catupres the variability for aggregated tick-by-tick prices. It actully just takes the first and last tick-by-tick prices into account.




Figure 2.13 The absolute daily returns and GARCH (1,1) volatility forecasts

In contrast, Figure 2.14 suggests the daily GARCH forecasts are much more in line with the ex post daily return variability measure $\sqrt{\sum_{n=1}^{N_t} r_{t,n}^2}$ and $\sum_{n=1}^{N_t} |r_{t,n}|$. These two measures are calculated by tick-by-tick returns and account for the variability of every tick-by-tick return.




Figure 2.14 The cumulative absolute returns, realized volatility and GARCH (1,1) forecasts

To further illustrate this issue, I examine the correlations between the daily GARCH forecasts and the other three series. The results are reported in Table 2.7.

	$\sum_{n=1}^{N_t} r_{t,n} $	P value	$ R_t = \left \sum_{n=1}^{N_t} r_{t,n} \right $	P value	$\sqrt{\sum_{n=1}^{N_t} r_{t,n}^2}$	P value
σ_t	0.553814319	1.17E-21	0.236701	1.49E-04	0.559359	3.80E-22
	$[\sum_{n=1}^{N_t} r_{t,n}]^2$	P value	R_t^2	P value	$\sum_{n=1}^{N_t} r_{t,n}^2$	P value
σ_t^2	0.54229162	1.14E-20	0.230797	2.19E-04	0.548984	3.08E-21

Table 2.7 Correlations between GARCH forecasts and other ex post daily variability measures

Table 2.7 confirms the finding from Figures 2.13 and 2.14. The correlation between the GARCH forecasts and the realized volatility is around 0.559. That is, 31.25 percent of the variation in the realized volatility can be explained by the MA (1)-GARCH (1,1) model. Besides, the correlation between the GARCH forecast and cumulative absolute returns is also as high as 0.5538. In contrast, the correlation between the GARCH forecast and absolute daily returns is only 0.2367. Thus, only about 6 percent of the variation of the daily absolute returns can be captured by the MA (1)-GARCH (1,1) forecasts.

Moreover, since the GARCH forecasts are based on daily observations, the strong correlation between GARCH forecasts and realized variance cannot be explained by only intraday volatility pattern. No matter what the intraday volatility pattern is, it should be annihilated when returns are measured at daily frequency. Consequently, it suggests that

the ARCH effects observed at daily frequency have strong impact on the intraday volatility pattern. As a result, ignoring the pronounced ARCH effect at daily frequency and the corresponding GARCH-ARCH forecasts in intraday volatility modelling inevitably results in the loss of a large percentage of predictable intraday return variability.

According to the analysis and discussion in this section, a misspecification of the intraday volatility process might emerge if we cannot capture the daily ARCH effect. The empirical results of Figures 13, 14 and Table 2.7 justify the daily conditional component σ_t in (2.10) and support the use of the GARCH prediction to estimate the σ_t .

2.5.2 Intraday volatility periodicity

In this section, I present our estimation of the intraday volatility periodicity component $s_{t,n}$. In Andersen and Bollerslev (1998), the intraday returns are modelled by the following equation.

$$r_{t,n} = E(r_{t,n}) + \sigma_{t,n} s_{t,n} z_{t,n}, \quad (2.18)$$

where the conditional volatility component is $\sigma_{t,n}$, $s_{t,n}$ is the intraday periodicity component and $z_{t,n}$ is an *iid* sequence with zero mean and unit variance. For the estimation of $\sigma_{t,n}$, they use the estimator σ_t/\sqrt{N} where σ_t is the daily GARCH estimation. Therefore, the intraday periodicity component $s_{t,n}$ is considered as a stochastic process evolving with time and is estimated by flexible Fourier transform. Regarding the functional data analysis of intraday volatility periodicity are in Muller et al. (2007) and Andersen and Bollerslev (1998). Alternatively, Engle (2012) considers the $s_{t,n}$ as constants and grants complete freedom to the daily volatility component σ_t . Specifically, $s_{t,n} = s_{\tau,n} = s_n$ for all τ and t .

Here, I follow the methodology of Engle (2012) for the estimation of $s_{t,n}$ based on following considerations. First, if I assume that the $E(r_{t,n})$ in (2.18) to be zero and replace the intraday volatility component $\sigma_{t,n}$ with its estimator σ_t/\sqrt{N} . The equation now reduces to (2.1), which is

$$r_{t,n} = 1/\sqrt{N} \sigma_t s_{t,n} z_{t,n}.$$

The intraday component in (2.18) is now $s_{t,n}$. It evidently should be assumed as a stochastic process. In contrast, recall the specification of intraday returns in (2.10),

$$r_{t,n}/\sqrt{w_{t,n}} = \sigma_t \epsilon_{t,n} s_{t,n} z_{t,n},$$

where $\epsilon_{t,n}$ is the intraday volatility component and is stochastic. If I consider $s_{t,n}$ as a stochastic process, it is hard to assume that it should be independent of $\epsilon_{t,n}$ and further complicate the picture. In order to explicitly exhibit the intraday return dynamics, it seems natural to assume the periodicity component to be unconditional. Secondly, based on our discussion of realized variance in Chapter 1, I should have,

$$E(\sigma_t^2) = E\left(\sum_{n=1}^{N_t} r_{t,n}^2\right) = \sum_{n=1}^{N_t} E(w_{t,n} \sigma_t^2 s_{t,n}^2).$$

It intuitively suggests that $s_{t,n}^2$ should be considered as time invariant as well.

I now present the estimation of intraday periodicity component $s_{t,n}$. Since that $E\left(\frac{r_{t,n}^2}{\sigma_t^2}\right) = E(w_{t,n} \epsilon_{t,n}^2 s_{t,n}^2 z_{t,n}^2) = w_{t,n} s_{t,n}^2$, under the assumption that $w_{t,n} = w_{\tau,n}$ for all t and τ (time-aggregated returns), a natural way to estimate the $s_{t,n}^2$ is

$$\widehat{s_{t,n}^2} = \frac{1}{T} \sum_{t=1}^T \frac{r_{t,n}^2}{w_{t,n} \sigma_t^2}. \quad (2.19)$$

This is reasonable for intraday return series formed from identical time interval (time-aggregated returns). Since my tick-by-tick and transaction-aggregated are not equally spaced in time, in general the n th return of day t , $r_{t,n}$ and n th return of day τ , $r_{\tau,n}$ are not sampled from the same time interval if $t \neq \tau$. I circumvent the difficulty by followings.

Let $x_{t,n}$ represent the time of the n th transaction at day t . Since $r_{t,n} = \ln\left(\frac{P_{t,n+1}}{P_{t,n}}\right)$, it can be viewed as the return of day t on period from $x_{t,n}$ to $x_{t,n+1}$. I now consider the intraday volatility pattern during the period $x_{t,n}$ to $x_{t,n+1}$ on each trading day. For another day k , let $k(x_{t,n+1}) = \max\{s | x_{k,s} \leq x_{t,n+1}, s = 1, 2, 3, \dots\}$, $k(x_{t,n}) = \min\{s | x_{k,s} \geq x_{t,n}, s = 1, 2, 3, \dots\}$. Let $S_k(t, n) = \{s | k(x_{t,n}) \leq s \leq k(x_{t,n+1})\}$. The prices at $x_{t,n}$ and $x_{t,n+1}$ on day k are conjectured by taking the average of adjacent two transactions in case that $S_k(t, n) = \emptyset$.

I then define $V_k(t, n) = \frac{1}{S[S_k(t, n)]} \sum_{s \in S_k(t, n)} r_{k,s}^2 / w_{k,s} \sigma_k^2$ where $S[S_k(t, n)]$ is the number of elements in $[S_k(t, n)]$. Further, $\widehat{s_{t,n}^2} = \frac{1}{N} \sum_{k=1}^N V_k(t, n)$ is applied to estimate the intraday volatility periodicity pattern during that period. The Matlab code for the estimation is given in Appendix C.

My intraday periodicity estimation procedure allows the shape of the periodic pattern depend on the estimated trading period. It turns out to be important for the tick-by-tick return series and transaction-aggregated return series for which the observations are not equally spaced in time. Besides, my approach uses the full tick-by-tick return series to accurately capture the periodicity.

However, due to the extremely large tick-by-tick sample size, the searching algorithm used to find out $S_k(t, n) = \{s | k(x_{t,n}) \leq s \leq k(x_{t,n+1})\}$ requires very long execution time. Moreover, the random-access memory requirement for the computation is beyond the capacity of my equipment. However, it is possible for traders with high technique capacity to use the tick-by-tick based algorithm to estimate the intraday periodic volatility pattern. Thus, I choose the time-aggregated returns to examine our specification since searching a regular partition of time massively reduce the computational difficulties.

2.5.3 Econometrics regarding estimation

In this section, I discuss statistical properties regarding my model estimation. I estimate the daily conditional volatility component σ_t use the MA (1)-GARCH (1,1) model first. Although in principle the parameters of the model in (2.10) should be estimated simultaneously, I do not adapt this standard procedure according to the following considerations. First, the daily GARCH model estimation generally requires a relatively large sample. However, the corresponding tick-by-tick high frequency data are not available. Second, a tick-by-tick transaction records of several years will have huge number of observations. E.g. My tick-by-tick transaction records of symbol SPY during 01/02/2014-12/31/2014 has 79156264 observations. Consequently, it rises the practical

difficulty of storing and computing data of such magnitude in a joint estimation. Finally, since σ_t is estimated by MA (1)-GARCH (1,1) without a clear specification in (2.10), it gives the flexibility to use other daily volatility measures.

Thus, the estimation now is a two-step procedure. The first step I estimate the intraday volatility periodicity component as discussed in Section 2.5.2. The second step I estimate the GARCH (1,1) model for the stochastic intraday volatility component $\epsilon_{t,n}$. In general, results from a multi-step estimation procedure can be flawed by the errors that are induced by errors in the previous estimation steps. In order to show the estimation are consistent, I apply the generalized methods of moments methodology (GMM).

Let ψ be the k_1 parameters estimated at the first step and θ be the k_2 parameters estimated at the second step. Besides, suppose $g_1(\psi)$ and $g_2(\psi, \theta)$ are the k_1 and k_2 moments conditions which specify the parameters by true moments. Denote the corresponding sample sum as $g_{1,M}$ and $g_{2,M}$. Let $g_M = (g'_{1,M}, g'_{2,M})'$, $\delta = \begin{pmatrix} \psi \\ \theta \end{pmatrix}$ and $g(\delta) = \begin{pmatrix} g_1(\psi) \\ g_2(\psi, \theta) \end{pmatrix}$. The GMM estimator of the parameter can be presented as

$$\begin{pmatrix} \hat{\psi} \\ \hat{\theta} \end{pmatrix} = \arg \min g_M' I g_M = \arg \min g_M' g_M$$

I is the diagonal matrix. In order to solve the above equation, ψ must solve the equation in first row and θ should solve the equation in the second row given $\hat{\psi}$. According to Newey and McFadden (1994), if $\hat{\psi}$ and $\hat{\theta}$ are consistent estimators of the ψ and θ , under quite general conditions (Newey and McFadden 1994) of g_M , the following estimator $\sqrt{M} \begin{pmatrix} \hat{\psi} & \psi \\ \hat{\theta} & \theta \end{pmatrix}$ is consistent and asymptotically normal. Specifically, it converges in distribution to the normal distribution $N(0, G^{-1} \Lambda G^{-1'})$ where $G = E(\partial g(\delta) / \partial \delta)$ and $\Lambda = E(g(\delta) g(\delta)')$.

As showed by Hansen (1982) and Engle (2012), the matrix $\sqrt{M} \begin{pmatrix} \hat{\psi} & \psi \\ \hat{\theta} & \theta \end{pmatrix}$ can be consistently estimated by using sample averages to replace expectations and using estimation of parameters to replace parameters.

In the context of my estimation, let $r_{t,n}$, $\epsilon_{t,n}$, σ_t and $s_{t,n}$ be defined as in equation 2.10. I have

$$\begin{aligned}
g_{1,M}(\psi) &= g_{1,M} = \begin{pmatrix} 1/T \sum_{t=1}^T \frac{r_{t,1}^2}{\sigma_t^2} - s_{t,1}^2 \\ \vdots \\ 1/T \sum_{t=1}^T \frac{r_{t,N_t}^2}{\sigma_t^2} - s_{t,N_t}^2 \end{pmatrix}, \\
g_{2,M}(\hat{\psi}, \theta) &= g_{2,M} = 1/T \sum_{t=1}^T \sum_{n=1}^{N_t} \nabla_{\theta} (\log(\epsilon_{t,n}^2) + \left(\frac{r_{t,1}^2}{\sigma_t^2 s_{t,n}^2 \epsilon_{t,n}^2} \right)), \\
G &= \frac{1}{M} \sum \begin{pmatrix} \nabla_{\psi} g_1 & 0 \\ \nabla_{\theta} g_2 & \nabla_{\theta} g_2 \end{pmatrix}, \\
\frac{1}{M} \sum \begin{pmatrix} g_{1,t}^2 & g_{1,t} g_{2,t} \\ g_{1,t} g_{2,t} & g_{2,t}^2 \end{pmatrix} &\rightarrow \Lambda. \tag{2.20}
\end{aligned}$$

The convergence is in the sense of convergence in probability measure. Moreover, the $\hat{\psi}$ and $\hat{\theta}$ estimated at the first and second stage are consistent estimator for ψ and θ clearly. Under the assumption of stationarity and ergodicity, Large Number Theorem verifies the consistency of $\hat{\psi}$. Consistency of the GARCH parameters can be found in Lee and Hansen (1994) and Bollerslev (1992).

2.6 Empirical results

In this section I present the empirical results of my intraday volatility specification of (2.10). I evaluate the specification from mainly two perspectives. First, I analyze the intraday returns normalized by one-day-ahead GARCH forecasts, the intraday returns normalized by one-day-ahead GARCH forecasts and intraday volatility periodicity pattern.

I use the terminology “*normalized returns*” to refer to the intraday returns normalized by one-day-ahead GARCH forecasts. Specifically, $\widetilde{r}_{t,n} = r_{t,n}/\widehat{\sigma}_t$. The terminology “*filtered returns*” are used to describe intraday returns normalized by one-day-ahead GARCH forecasts and intraday volatility periodicity pattern. Specifically, $\widehat{r}_{t,n} = r_{t,n}/w_{t,n}\widehat{\sigma}_t\widehat{s}_{t,n}$. As explained in previous section 2.5.2, the de-seasonal technique requires large amount of computation for the transaction-aggregated returns since it is based on searching tick-by-tick transactions records. Thus, I only exhibit the filtered returns for time-aggregated returns. Since those returns are sampled from identical time intervals, the duration $w_{t,n}$ in

(2.10) can be defined as unit.

Second, I present the model estimation for time-aggregated returns at various frequencies. I compare the estimated parameters with the theoretical aggregational prediction of GARCH (1,1) from Nelson (1992) and analyze the model residuals.

2.6.1 Normalized intraday returns and filtered intraday returns

In this section I present the analysis regarding normalized intraday returns and filtered intraday returns. I divide the analysis into two sections. Section 2.6.1.1 discusses the analysis of the normalized intraday returns. Section 2.6.1.2 represents the analysis of the filtered intraday returns.

2.6.1.1 Normalized intraday returns

Figure 2.15 presents the one-day-ahead MA (1)-GARCH (1,1) estimates for the period during 01/02/2014-12/31/2014. The details of the estimation are presented in section 2.5.1. I now present a summary of descriptive statistics for the normalized intraday returns $r_{t,n}/\hat{\sigma}_t$.

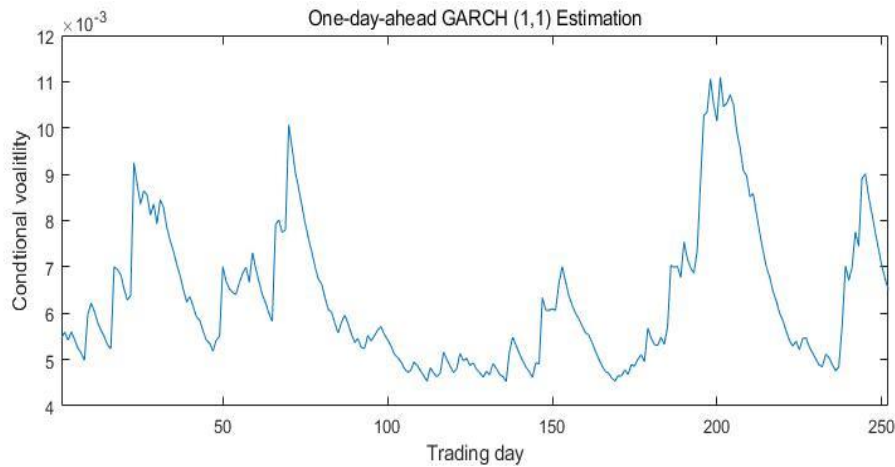


Figure 2.15 One-day-ahead MA (1) -GARCH (1,1) estimation

	Sample Size	Mean	Variance	Kurtosis	Skewness	Max	Min	p_1	$(a)p_1$	$(s)p_1$
30Second	196560	1.25E-07	0.001204	167.665	-0.13508	1.385718	-1.46815	-0.10463	0.312765	0.489134
1Minute	98280	2.51E-07	0.002231	64.42646	-0.1912	1.408863	-1.45948	-0.08058	0.282229	0.467608
1.5Minute	65520	3.76E-07	0.00321	48.6927	-0.16571	1.361122	-1.49995	-0.06177	0.268761	0.453377
3Minute	32760	7.52E-07	0.006195	20.9807	-0.26811	1.245272	-1.63514	-0.04921	0.240136	0.378386
5Minute	19656	1.25E-06	0.0095	6.321807	-0.14592	0.731673	-0.83181	-0.01585	0.19892	0.15415
10Minute	9828	2.51E-06	0.018818	5.993268	-0.22319	0.814763	-0.98106	-0.01946	0.190533	0.125991
15Minute	6552	3.76E-06	0.027404	5.769621	-0.23229	0.945254	-1.09734	0.003031	0.181626	0.105811
30Minute	3276	7.52E-06	0.055366	5.137176	-0.25934	1.091958	-1.15388	0.008075	0.152081	0.107696
78Minute	1260	1.96E-05	0.146711	4.904028	-0.39886	1.448271	-1.88065	0.068372	0.167124	0.112384

Table 2.8 The normalized time-aggregated returns

	Sample Size	Mean	Variance	Kurtosis	Skewness	Max	Min	p_1	$(a)p_1$	$(s)p_1$
T=400	198261	-1.9E-06	0.001155	1308.24	0.189044	3.307417	-3.29112	-0.0962854	0.253815	0.499309
T=800	98942	-6.1E-06	0.002064	210.4466	0.008546	2.210617	-2.19975	-0.047105	0.182147	0.495312
T=1200	65893	-9.5E-06	0.003171	372.8296	0.021205	3.31828	-3.31285	-0.0569775	0.194002	0.497701
T=2400	32876	0.000106	0.005926	113.042	2.467613	3.35631	-0.55958	0.00530582	0.064171	0.001178
T=4000	19663	1.99E-05	0.009473	3.529551	-0.04754	0.557197	-0.58379	0.00794604	0.07018	0.088125
T=8000	9763	-2.2E-05	0.019155	3.4	-0.05032	0.755934	-0.53138	-0.0011895	0.052102	0.050723
T=12000	6468	0.000195	0.028742	3.314235	-0.15596	0.696386	-0.64206	-0.0049805	0.059878	0.055918
T=24000	3162	0.000689	0.057648	3.233583	-0.15545	0.891024	-0.91855	-0.0108197	0.032436	0.032172
T=62000	1165	-0.00216	0.161005	3.056574	-0.13479	1.343841	-1.22803	0.01098296	0.082668	0.060503

Table 2.9 The normalized transaction-aggregated returns

Tables 2.8 and 2.9 give the descriptive statistics of the time-aggregated returns and transaction-aggregated at various frequencies, where p_1 , $(a)p_1$ and $(s)p_1$ are the first lag autocorrelations of the normalized returns, absolute normalized returns and squared normalized returns respectively.

Although I again detect little interests in the sample mean, the skewness and kurtosis have intriguing features. Regarding the kurtosis, it turns out that when returns are measured at the highest frequencies, the normalized time-aggregated returns have much smaller kurtosis compared with the transaction-aggregated returns. For instance, the kurtosis of the normalized 30-second, 1-minute, 1.5-minute and 3-minute returns series are 167.67,

64.43, 48.69 and 20.9807 respectively. Meanwhile, the kurtosis of the corresponding normalized T-400, T-800, T-1200 and T-2400 transaction aggregated returns series are 1308.24, 210.45, 372.83 and 113.04 respectively. In other words, transaction-aggregated returns are more fat-tailed compared with time-aggregated returns at the highest frequencies. However, when returns are measured at relatively low level of intraday frequency, the normalized transaction-aggregated returns exhibit much regular kurtosis compared with corresponding normalized time-aggregated returns. For example, the kurtosis of the corresponding normalized T-4000, T-8000, T-12000 and T-24000 transaction aggregated returns series are 3.53, 3.4, 3.33 and 3.23 respectively. However, the corresponding normalized 5-minute, 10-minute, 15-minute and 30-minute return series have kurtosis of 6.32, 5.99, 5.77 and 5.13 respectively. This finding is consistent with our discussion regarding the two different sampling methods in Chapter 1. Observe the sharp drop of kurtosis between the T-2400 and T-4000 normalized transaction aggregated returns.

Regarding the skewness, note that the normalized T-4000 and T-8000 return series have skewness that are very close to 0. In contrast, the skewness of normalized time-aggregated returns varies and present no obvious dependence on the return frequencies.

Another important empirical finding comes from the last three columns in Tables 2.8 and 2.9. First, both normalized intraday time-aggregated and transaction-aggregated returns are approximately uncorrelated, although there is a small, but significant, negative first order autocorrelations at the highest frequencies. This is explained the bid-ask bounce as mentioned previously.

Second, in general the normalized intraday returns still exhibit significant volatility clustering feature. For instance, the normalized 1-minute, 1.5-minute and 3-minute squared returns have the first lag autocorrelations of 0.4676, 0.4533 and 0.3784 respectively. Meanwhile, the first lag autocorrelations of these three normalized time-aggregated returns are -0.0806, -0.0618 and -0.0492 that are neglectable from the economic perspective. The normalized transaction-aggregated returns exhibit similar characteristics. For example, the normalized T-800 and T-1200 squared returns have very strong positive first lag

autocorrelations of 0.4953 and 0.4978 and meanwhile the first lag autocorrelation of the normalized T-800 and T-1200 return series are -0.0481 and -0.0570 . This finding suggests the intraday returns exhibit strong ARCH effect even after normalizing by ARCH-GARCH daily volatility forecasts. This implicitly supports our specification of (2.10). However, this intraday ARCH effect generally dies out when return frequencies decrease as suggested by the monotonically decreasing values of the last two columns in Tables 2.8 and 2.9.

Third, compared with the normalized time-aggregated return series for which the intraday ARCH effect is significant even at very low intraday frequency, the ARCH effect of normalized transaction-aggregated return series vanishes very quickly. For instance, when the number used to aggregate the tick-by-tick transactions exceeds 2400, the first lag autocorrelation of normalized squared transaction-aggregated returns and normalized absolute transaction-aggregated are of very small values that can be assumed zero plausibly.

Under (2.10), The normalized time-aggregated intraday returns $\frac{r_{t,n}}{\sigma_t} = \epsilon_{t,n} S_{t,n} Z_{t,n}$. Thus, in addition to the discussed conditional heteroskedasticity, it should present the intraday periodicity pattern as well. I here present the correlogram of the normalized 5-minute absolute return series and the normalized 10-minute absolute return series for up to five trading days.




Figure 2.16 The correlogram of normalized 5-minute absolute returns




Figure 2.17 The correlogram of normalized 10-minute absolute returns

Figures 2.16 and 2.17 exhibit strong periodic pattern as expected. Compared with the Figures 2.5 and 2.6 in section 2.3, a salient feature is that the normalized intraday returns now no longer constantly possess significant autocorrelations for the lags up to five trading days.

2.6.1.2 The filtered intraday returns

In this section, I discuss the filtered intraday returns $\widehat{r}_{t,n} = r_{t,n}/\widehat{\sigma}_t \widehat{s}_{t,n}$. I first present a graphic illustration of our estimation of $s_{t,n}$. To be consistent with the previous description, I use the 5-minute time intervals to present the intraday periodicity.

Figure 2.18 presents the expected U-shape of intraday periodicity. It starts at the highest level when market opens and then gradually decreases with time. It reaches the lowest value at the forty 5-minute trading interval that is around 12:50. After that, it increases with time until the market is closed.




Figure 2.18 The estimated intraday periodicity for each 5-minute time interval

	Sample Size	Mean	Variance	Kurtosis	Skewness	Max	Min	p_1	$(a)p_1$	$(s)p_1$
30Second	196560	-0.00037	1.000005	16.04269	-0.04466	15.14489	-15.1776	-0.04421	0.237787	0.39142
1Minute	98280	-0.00047	1.00001	12.22964	-0.08725	14.24929	-14.3882	-0.05103	0.228836	0.346531
1.5Minute	65520	-0.00073	1.000015	12.90129	-0.08684	13.81001	-13.7784	-0.0321	0.217655	0.354665
3Minute	32760	-0.00078	1.00003	10.25224	-0.09803	12.36469	-12.5788	-0.03602	0.205242	0.321281
5Minute	19656	-0.00078	1.00005	5.63634	-0.13308	6.100498	-8.11604	-0.02155	0.171075	0.145198
10Minute	9828	-0.00118	1.0001	5.424702	-0.18769	5.90282	-6.74803	-0.02201	0.169528	0.125574
15Minute	6552	-0.00212	1.000148	5.408844	-0.24594	5.388117	-7.40249	0.009641	0.160309	0.099098
30Minute	3276	-0.00411	1.000288	4.959931	-0.27933	4.463494	-5.19666	0.012625	0.143373	0.108887
78Minute	1260	-0.00165	1.000792	4.869385	-0.34463	4.633105	-4.41664	0.069233	0.181823	0.128609

Table 2.10 The filtered time-aggregated returns

Table 2.10 summaries the descriptive statistics regarding the filtered time-aggregated returns at different frequencies. Observe the variance now is almost unity for filtered time-aggregated returns as specified under (2.10). The filtered time-aggregated returns still present excessive high kurtosis and negative skewness, suggesting that the unconditional distribution is not normal. An interesting finding is that the filtered time-aggregated returns

now present little evidence for autocorrelations. Specifically, the first lag autocorrelation for the 30-second, 1-minute, 1.5-minute, 3-minute and 5-minute return series are now - 0.044, -0.051, -0.032, -0.036 and -0.022 respectively. These values are still significant but of small magnitude. In contrast, the squared filtered time-aggregated return series still present strong autocorrelations as suggested by the last column of Table 2.10. This implies the existence of conditional heteroskedasticity.

I carry out the ARCH test (Engle 1988) to verify the conditional heteroskedasticity as in Section 2.5.1. Specifically, we test the null hypothesis that there is no ARCH effect in the filtered time-aggregated return series against the alternative hypothesis of an ARCH model with two lagged squared innovations that is equivalent to a GARCH (1,1) model locally.

	P values	Value of TR^2	Critical value of TR^2
30Second	0	32328.71024	9.210340372
1Minute	0	12087.05103	9.210340372
1.5Minute	0	8450.709077	9.210340372
3Minute	0	3402.702996	9.210340372
5Minute	0	728.5254549	9.210340372
10Minute	0	303.1955374	9.210340372
15Minute	0	123.6137445	9.210340372
30Minute	0	113.2577976	9.210340372
78Minute	0	86.10606377	9.210340372

Table 2.11 ARCH test of filtered time-aggregated returns

As the suggested by Table 2.11, the null hypothesis of no ARCH effect is rejected at all frequencies. The significance level is 0.01. In other words, the time-aggregated returns after normalization by daily conditional variance and intraday periodicity still present strong conditional heteroskedasticity. Finally, I present the correlogram of the absolute filtered 5-minute and 10-minute time-aggregated return series for up to five days.




Figure 2.19 The correlogram of filtered 5-minute absolute returns

The autocorrelations of filtered 5-minute absolute returns decreases now very slowly through the lags since the intraday periodic pattern has been excluded by $\widehat{r}_{t,n} = r_{t,n} / \widehat{\sigma}_t \widehat{s}_{t,n}$. This is in line with autocorrelation structure of GARCH model. The U-shape in the Figures 2.15 and 2.16 now no longer exists. It suggests that my estimation of the intraday volatility periodicity is effective. The following Figure 2.20 is equally telling.




Figure 2.20 The correlogram of filtered 10-minute absolute returns

2.6.2 Results of GARCH estimation and volatility persistence

In this section I present the results of the GARCH estimation. I first exhibit the inadequacy of the standard GARCH modelling for describing intraday volatility process. I then present our GARCH estimation of the filtered time-aggregated return volatility. I focus on the comparison between the estimates and the theoretical aggregation results from Nelson (1990,1992). Finally, I present the model residuals using the example of 5-minute and 10-minute filtered returns.

There are vast literatures regarding the modelling of high frequency return dynamics using GARCH (1,1) model. Thus, I first examine the effectiveness of the GARCH (1,1) in modelling raw intraday time-aggregated returns. In order to capture the significant first order autocorrelation, I choose the MA (1)-GARCH (1,1) specification.

$$\begin{aligned} R_{t,n} &= u + \theta \varepsilon_{t,n-1} + \varepsilon_{t,n}, \\ \varepsilon_{t,n} &= \sigma_{t,n} z_{t,n}, \\ \sigma_{t,n}^2 &= \omega + \alpha \varepsilon_{t,n-1}^2 + \beta \sigma_{t,n-1}^2, \end{aligned} \quad (2.21)$$

where $E_{t,n-1}[(\varepsilon_{t,n}^2)] = \sigma_{t,n}^2$ is the conditional variance and $z_{t,n} \sim N(0,1)$ and is *iid*.

The estimation is based on the quasi-maximum likelihood methods. Given the intraday conditional heteroskedasticity discussed previous, it seems natural to choose the GARCH (1,1) specification for the intraday variance modelling since it usually represents a reasonable approximation. Besides, estimating the GARCH (1,1) model across all time-aggregated return frequencies yields meaningful comparisons between estimated parameters. In order to describe the volatility persistence implied by the GARCH (1,1) parameters, I use the “half-life” that is the number of periods the process takes for half of the expected reversion backs to the unconditional variance. Specifically, $-\ln 2 / \ln(\alpha + \beta)$. The “mean lag” is also supplied as an additional measure of the volatility persistence. i.e. $\alpha(1 - \alpha - \beta + \alpha\beta)^{-1}$.

Nelson (1990,1992) and Drost and Nijman (1993) provide a framework to assess the GARCH parameter estimates at various sampling frequencies. Specifically, consider the GARCH (1,1) at the daily frequency. Let $\alpha^{(1)}$ and $\beta^{(1)}$ be the parameters for the GARCH

(1,1) daily return model, then the $\alpha^{(t)}$ and $\beta^{(t)}$ that are parameters for the GARCH model for the t -day returns should suffice the equation,

$$\alpha^{(t)} + \beta^{(t)} = (\alpha^{(1)} + \beta^{(1)})^t. \quad (2.22)$$

Equation (2.22) immediately implies that the estimated t -day GARCH half-life should be related with daily GARCH half-life as the following equation,

$$-\ln 2 / \ln (\beta^{(t)} + \alpha^{(t)}) = -\ln 2 \frac{1}{\ln(\alpha^{(1)} + \beta^{(1)})^t} = -\frac{1}{t} \frac{\ln 2}{\ln(\alpha^{(1)} + \beta^{(1)})}. \quad (2.23)$$

As a result, the estimated half-life of GARCH models of return series with different sampling frequencies should be stable if they are normalized by corresponding frequencies. Besides, as showed by Nelson (1990, 1992), the relations hold under general conditions of GARCH (1,1), no matter whether the model is mis-specified at some frequencies or not. Empirical evidence for the relations between daily and lower frequency GARCH estimation can be found in Drost and Nijman (1993), Drost and Werker (1996) and Bollerslev (1997).

	α	Standard error	β	Standard error	$\alpha + \beta$	Half-life	Median lag
30-Second	0.157	0.0065542	0.918	0.0022418	1.075	∞	∞
1-Minute	0.139047	0.0015039	0.681	0.0051612	0.820047	3.493795158	0.506280349
1.5-Minute	0.114605	0.0011804	0.434	0.0053909	0.548605	1.73177985	0.343036541
3-Minute	0.290064	0.0067264	0.195893	0.0029211	0.485957	2.881572988	1.524343333
5-Minute	0.335584	0.0106232	0.250198	0.0068206	0.585782	6.480334013	3.368093719
10-Minute	0.318509	0.0143567	0.472546	0.0092005	0.791055	29.57271835	8.860894976
15-Minute	0.278234	0.0100601	0.575654	0.014528	0.853889	65.82402449	13.62654377
30-Minute	0.160019	0.0110557	0.7568	0.0083571	0.916819	239.4433497	23.49961427
78-Minute	0.148918	0.0177689	0.791006	0.0153911	0.939924	872.6371124	65.30338637

Table 2.12 GARCH estimation of intraday time-aggregated returns

The half-life and median-lag in Table 2.12 are converted into unit minute. Consider the daily MA (1)-GARCH (1,1) estimation that gives the $\alpha = 0.104$ and $\beta = 0.873$, the results of Drost and Nijman (1993) now suggest that the intraday returns should follow weak GARCH (1, 1) processes with $\alpha + \beta$ approaching 1 and α approaching 0 as the frequencies increase. However, the α and β in Table 2.12 behave erratically since the sum of the two parameters does not exhibit such tendency clearly. Moreover, the half-life and median-lag present dramatic variations along with the sample frequencies. It is very clear

that a direct use of GARCH (1,1) specification of intraday return volatility is seriously in doubt. I now present the results of GARCH modelling of the filtered time-aggregated return $R_{t,n}/\widehat{\sigma}_t\widehat{s}_{t,n}$ in further details.

	ω	Stand error	T statistics	α	Standard error	T statistics	β	Standard error	T statistics	$\alpha + \beta$	Half- life	Median lag
30second	0.00471	6.04E-05	77.95481167	0.04074	0.000236967	171.939	0.95588	0.000240364	3976.801	0.9966	102.422	0.4813303
1minute	0.00575	0.000113	50.86504525	0.04562	0.000533077	85.58773	0.94968	0.00054553	1740.836	0.9953	147.25	0.9500191
1.5minute	0.00592	0.0001692	34.97214726	0.05346	0.000738496	70.29338	0.94265	0.000760529	1276.446	0.9961	266.826	1.4772541
3minute	0.0087	0.0003331	26.11107358	0.05715	0.001174732	48.64809	0.93564	0.001247306	750.1245	0.9928	287.116	2.8251007
5minute	0.00936	0.0005619	16.65479684	0.06235	0.002083684	29.92373	0.9295	0.002230022	416.8098	0.9918	423.339	4.7158355
10minute	0.01609	0.0013179	12.20932696	0.08228	0.003900577	21.09331	0.90385	0.00415363	217.6043	0.9861	496.054	9.3240088
15minute	0.02301	0.002541	9.05468895	0.111	0.006779362	16.37279	0.87125	0.007439406	117.1129	0.9822	580.461	14.546339
30minute	0.03857	0.0043299	11.19373674	0.10985	0.009813608	11.19374	0.85459	0.01109267	77.04099	0.9644	574.33	25.46061
78minute	0.11523	0.0195757	5.886293674	0.18905	0.028076156	6.733532	0.70614	0.033930637	20.81114	0.8952	488.296	61.877751

Table 2.13 GARCH estimation of filtered intraday time-aggregated returns

Table 2.13 suggests that the estimation now is much more in line with theoretical predictions. First, the half-life estimated from various frequencies are relatively stable now. For instance, the half-life estimated from the filtered 5-minute, 10-minute, 30-minute and 78-minute return series are 423.34, 496.05, 580.46, 574.33 and 488.30 respectively. Second, the parameters α and β are strikingly regular as the theoretical prediction. Observe now α monotonically decreases and meanwhile β monotonically increases as the sample frequencies increase. Since the daily MA (1)-GARCH (1,1) estimation that gives the $\alpha = 0.104$ and $\beta = 0.873$, according to (2.22), I should observe that $\alpha + \beta$ approaches one and α approaches zero as the frequencies increases, which is perfectly matched by the empirical results.

I now present the residual correlogram of the filtered 5-minute return series and filtered 10-minute return series.




Figure 2.21 Residuals correlogram of the filtered 5-minute returns




Figure 2.22 Residuals correlogram of the filtered 10-minute returns

Figures 2.21 and 2.22 suggest that the intraday conditional heteroskedasticity in Figures 2.19, 2.20 and Table 2.11 is successfully captured by the GARCH (1,1) model. Figure 2.23

gives the residual Q-Q plot of the filtered 5-minute returns, the residuals still present the fat-tailed property that is widely identified by high frequency financial research.




Figure 2.23 The Q-Q plot of filtered 5-minute residuals

2.7 Conclusion

In this chapter I analyze the intraday volatility pattern of the SPY on NYSE. I show that standard time series modelling techniques might draw erroneous conclusions due to the distortion induced by strong intraday volatility periodicity. In order to draw meaningful results from the high frequency intraday data, the modelling of the intraday periodicity pattern is necessary. Based on the tick-by-tick transaction records, I develop a non-parametric intraday periodicity modelling methodology that can be used in the data set where observations are not equally spaced in time. However, the estimation procedure for intraday volatility periodicity requires large computation capacities and durations are given exogenously. Further research could be carried out to jointly model the durations and intraday volatility. I also examine the relation between the daily and intraday conditional heteroskedasticity. I find out the intraday conditional heteroskedasticity cannot be fully explained by the long-memory characteristics of daily observations. I prove that the combination of intraday volatility periodicity and intraday conditional heteroskedasticity helps to explain the intraday volatility pattern.

In the analysis of the intraday return dynamics, I find out that a key element that relates closely to the estimation of periodicity and the modelling of instantaneous volatility is the duration between consecutive transactions that I will elaborate in the Chapter 3.

Chapter 3

Autoregressive Conditional Duration modelling of durations

3.1 Introduction

The rising availability of tick-by-tick transaction data in financial markets provides the possibility to directly investigate the characteristics associated with each tick transaction record. There are three main characteristics of each tick transaction record: execution price, trading volumes and execution time. A distinctive feature of the tick-by-tick transaction data is that the observations are no longer equally spaced in time. This feature challenges standard time series models such as the ARCH-GARCH family since such models are based on observations equally spaced in time. Early high frequency research applied the time-aggregated return series to circumvent this challenge (Hafner 1996, Guillaume, Pictet and Dacorogna 1995 and Olsen and Pictet 1997).

However, this procedure faces two main difficulties in practice. First, it inevitably causes the loss of the information contained in the order of the transaction flows. As argued by market microstructure research, timing of the transactions could be vital to the understanding of economics behind the price behavior. For instance, the duration between transactions could be viewed as a signal of asymmetric information among traders (Easley and O'Hara 1992, 1995 Diamond and Verrecchia 1987,1991 and Goodhart and O'Hara 1997). Thus, a coherent description of the tick-by-tick prices dynamics should base on the modelling of durations. Second, the intraday volatility periodicity discussed in Chapter 2 makes the choice of the optimal time interval used to aggregate tick-by-tick transactions very difficult. Specifically, if a short interval is chosen, there could be many intervals that contain no new information. Consequently, a specific form of heteroskedasticity corresponding to this short interval will be induced into the time-aggregated return sample. In contrast, the use of a long interval will reduce the sample size significantly. The discussion of the characteristics of time-aggregated returns can be found in Chapter 1.

In order to give an explicit parametric description of the “waiting-time” that is the duration between consecutive transactions, Engle and Russell (1998) proposed the Autoregressive Conditional Duration (ACD) model based on the point process. A detailed introduction of the point process is presented in Daley and Vere-Jones (2006). The ACD model can be viewed as an analogy of the GARCH model for the duration modelling. Since its introduction, the ACD model has been widely developed and accepted as a tool in the modelling of financial data that are irregularly spaced in time. Many theoretical efforts have been made on the extension of the standard ACD model. See, e.g., Lunde (1996), Bauwens and Giot (2000), Engle (2000) and Zhang et al. (2001). Meanwhile, empirical research focused on the performance of the predictability of ACD models (Dufour and Engle 2000, Hautsch 2000 and Pacurar 2007).

There are two main questions related with the ACD models. The first one is: what is the theoretical relation among the various extensions of the ACD models and which one provides the most accurate modelling of the durations? The second one is: how can the ACD model be used to describe the prices dynamics?

In this chapter, I address the first question by testing different specifications of the ACD model. I focus on the comparison of different distributional assumptions. I use transaction-aggregated return series at various frequencies to examine different distributions. I also provide a review of both theoretical and empirical research with respect to the ACD models. I discuss the statistical properties of several ACD specifications along with some practical issues in applications. I find out that, compared with other distributions, Gamma distribution gives the best ACD modelling results for transaction-aggregated durations. Besides, the comprehensive introduction of ACD models could contribute to the understanding of the scope of research regarding durations modelling.

This chapter is organized as follows. Section 3.2 presents the discussion with respect to the ACD models. I analyze the theoretical framework of the ACD models and discuss application issues. Section 3.3 describes the data and some preliminary results with respect to the durations. I focus on the unconditional distributional properties of the tick-by-tick

durations. Section 3.4 presents the empirical result of the ACD modelling of duration dynamics. Section 3.5 is the conclusions.

3.2 ACD framework

In this section I analyze and present the framework of the ACD models. In order to present an exhaustive discussion, I first discuss the theoretical background of the ACD models in Section 3.2.1. Section 3.2.2 gives a discussion of the applications of ACD models.

3.2.1 Theoretical background of ACD

In this section I present the theoretical framework of the ACD models. Section 3.2.1.1 exhibits the point process that is the statistical background of the ACD models. I then discuss the ACD specification in Section 3.2.1.2 and analyze the relations between durations and other characteristics associated with each transaction such as volume and prices in Section 3.2.1.3.

3.2.1.1 Point Processes

Each tick-by-tick transaction record are associated with three basic characteristics: prices $p_{\tau,i}$, volumes $v_{\tau,i}$ and execution time $t_{\tau,i}$, where the index τ labels the trading day of the transaction and the index i indicates the order of the transaction during trading day τ . In order to explicitly present the point process, I now remove the day index τ and view the whole data set as a time series labeled just by the order index i .

As discussed previously, a salient feature of the tick-by-tick transaction records is that the durations between consecutive transactions are not equally spaced in time. However, they are ordered in time and therefore can be naturally considered as a point process. In probability theory, the point process is a collection of random variables that are randomly spaced in time. It is widely developed and applied in queueing theory and neuroscience. In 1998, Engle and Russell extended it to the high-frequency finance to model the trading process.

Specifically, consider the sequence of random variables $\{t_0, t_1, \dots, t_n, \dots\}$, where t_n represent the arrival time and $t_0 \leq t_1 \leq \dots \leq t_n \leq \dots$. In general, simultaneous occurrence of the events are allowed. In practice, “thinning” procedures usually are used to eliminate this possibility. Let the counting function $N(\tau)$ represents the total number of events that occur by time τ . $N(\tau)$ is a step function that has right limit and is left continuous at each point. The characteristics associated with each arrival time are considered as “marks”. In the context of high frequency financing, the trading volume and execution price of each transactions can be viewed as marks associated with the arrival time of transactions.

An important tool for the analysis of point process is the intensity function. With respect to the definition and detailed discussion of point process and intensity function, Daley and Vere-Jones (2006) and Cox and Isham (1980) present very comprehensive and exhaustive introductions. In this chapter, I focus on the point process in the context of high frequency financing. Specifically, the following specification of conditional intensity process is used to describe the point process,

$$\lambda(\tau | N(\tau), t_0, \dots, t_{N(\tau)}) = \lim_{h \rightarrow 0} \frac{\Pr(N(\tau+h) > N(\tau) | N(\tau), t_0, \dots, t_{N(\tau)})}{h}. \quad (3.1)$$

(3.1) implies that the current intensity is conditional the on past information of arrival time and the fact that there is no event since time $t_{N(\tau)}$. The conditional intensity process defined by (3.1) is also called the hazard function. The hazard function generally is defined as the ratio of the probability density function and survival function. It measures the instantaneous rate of the occurrence of events given that there are no events since the last previous event. Under (3.1), suppose f_i is the conditional probability density function for arrival time t_i . The log likelihood, L , of the whole sample is

$$L = \sum_{i=1}^{N(T)} \log f_i(t_i | t_0, \dots, t_{i-1}). \quad (3.2)$$

Observe that (3.2) can be rewritten in terms of the conditional intensity as

$$L = \sum_{i=1}^{N(T)} \log \lambda(t_i | i-1, t_0, \dots, t_{i-1}) - \int_{t_0}^T \lambda(s | N(s), t_0, \dots, t_{N(s)}) ds. \quad (3.3)$$

The equation (3.2) and equation (3.3) are used to for the maximum likelihood estimation of the model. It is very clear that the parametric specification of the conditional intensity process in (3.1) is critical to a successful point process modelling of high frequency trading pattern. Haustch (2000) provided a thorough review of the specifications of intensity

functions used in point process for financial modelling. From the perspective of probability theory, the conditional intensity fully describes the corresponding point process. E.g., Daley and Vere-Jones (2006) and Lancaster (1990). In this Chapter, I focus on the parameterization of (3.1) by Engle and Russell (1998) in the ACD framework.

3.2.1.2 ACD models

The ACD model parameterizes the point process by specifying the process of durations between consecutive events. Specifically, denote the i th duration between two events that occur at t_{i-1} and t_i as $w_i = t_i - t_{i-1}$. Thus, the sequence of durations now is $\{w_0, w_1, \dots, w_n, \dots\}$ and is nonnegative obviously. Let z_i be the marks such as volume and prices associated with the i th event. The durations and marks can be jointly presented as the sequence $\{(w_i, z_i)\}_{i=1,2,\dots,n,\dots}$. Thus, the joint conditional density of the i th element (w_i, z_i) can be presented as

$$(w_i, z_i) | F_{i-1} \sim f(w_i, z_i | w_1, w_2, \dots, w_{i-1}, z_1, z_2, \dots, z_{i-1}; \theta), \quad (3.4)$$

where F_{i-1} denotes the information set available at time t_{i-1} and $\theta \in \Theta$ is the set of parameters. Under the assumption of conditional independence, (3.4) can also be expressed as the product of marginal conditional density of the durations and the conditional density of the marks. Thus,

$$f(w_i, z_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta) = f_w(w_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_w) g(z_i | \bar{w}_i, \bar{z}_{i-1}; \theta_z), \quad (3.5)$$

where $\bar{w}_{i-1} = (w_1, w_2, \dots, w_{i-1})$ and $\bar{z}_{i-1} = (z_1, z_2, \dots, z_{i-1})$ are the sets of past information of the variables W and Z and f_w is the marginal conditional density function of the duration w_i with parameters θ_w and g is the conditional density function of the mark z_i with parameters θ_z . (3.5) describes the relation between durations and marks associated with events. It forms the theoretical foundation of modelling tick-by-tick prices dynamics under the ACD framework. For instance, Engle (2000) built the Ultra-High-Frequency GARCH model based on the ACD modelling of durations.

From (3.4) and (3.5), one can immediately write the log likelihood as

$$L(\theta_w, \theta_z) = \sum_{i=1}^{N(T)} [\ln f_w(w_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_w) + g(z_i | \bar{w}_i, \bar{z}_{i-1}; \theta_z)]. \quad (3.6)$$

Weak exogeneity refers to the case where the statistical estimation can be obtained by considering the conditional model (Engle, Hendry and Richard 1983 and Johansen 1991). Under the assumption that the durations are weakly exogenous to the marks, the right part and left part of (3.6) are usually maximized separately (Engle 2000). This property effectively simplifies the model estimation procedures. However, the weak exogeneity is usually imposed as an assumption since there is no widely accepted statistical test for it in the context of ACD.

With respect to the specification of conditional marginal density of the durations, Engle and Russell (1998) suggested the following specification. Denote the conditional expectation of the i th duration w_i as

$$E(w_i | \bar{w}_{i-1}, \bar{z}_{i-1}) \equiv \psi_i. \quad (3.7)$$

The ACD model assumes that

$$\psi_i = \psi(\bar{w}_{i-1}; \theta_\psi) \quad (3.8)$$

$$w_i = \psi_i \epsilon_i, \quad (3.9)$$

where ϵ_i is an iid sequence with density $p(\epsilon; \phi)$ and ψ is function of ψ_i with parameters θ_ψ . Observe that (3.9) implicitly requires that $E(\epsilon_i) = 1$. This assumption is without loss of generality since one can use $w_i = \psi_i \epsilon_i / E(\epsilon_i)$ equivalently. In estimation of the model, this assumption is usually released to allow for flexibility.

The specification of (3.8) requires the density function $p(\epsilon; \phi)$ to have a nonnegative support since the durations are nonnegative in nature. Besides, (3.8) suggests that $f_w(w_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_w) = f_w(w_i | \bar{w}_{i-1}; \theta_w)$. In other words, the conditional expectation of durations depends upon only the past durations. Consequently, the conditional density of the durations can be expressed as

$$f_w(w_i | \bar{w}_{i-1}; \theta_w) = f_w(w_i | \bar{w}_{i-1}; \theta_\psi, \phi),$$

$$f_w(w_i|\bar{w}_{i-1}; \theta_\psi, \phi) = \frac{1}{\psi_i} p\left(\frac{w_i}{\psi_i}; \phi\right). \quad (3.10)$$

Thus, the log likelihood function is,

$$\begin{aligned} L(\theta_\psi, \phi) &= \sum_{i=1}^{N(T)} \ln f_w(w_i|\bar{w}_{i-1}; \theta_\psi, \phi), \\ \sum_{i=1}^{N(T)} \ln f_w(w_i|\bar{w}_{i-1}; \theta_\psi, \phi) &= \sum_{i=1}^{N(T)} [\ln p\left(\frac{w_i}{\psi_i}; \phi\right) - \ln \psi_i]. \end{aligned} \quad (3.11)$$

Once the innovation term ϵ_i and the conditional expectation ψ_i are specified, then maximum likelihood estimates of θ_ψ and ϕ can be obtained by general numerical optimization procedures. Finally, denote the survival function associated with the density $p(\epsilon; \phi)$ of ϵ as S , the conditional intensity of the ACD model can be deduced from (3.1),

$$\lambda(t|N(\tau), t_0, \dots, t_{N(\tau)}) = \frac{p(\tau)}{S(\tau)} \left(\frac{\tau - t_{N(\tau)}}{\psi_{N(\tau)+1}} \right) \frac{1}{\psi_{N(\tau)+1}}. \quad (3.12)$$

The deduction details of equation (3.12) could be found in Engle and Russell (1998). The setup from (3.7) to (3.9) are very general and therefore allows a flexible structure of the ACD models. Specifically, a variety of ACD models can be obtained by combining different specifications for the conditional expectation ψ_i and the conditional density of ϵ_i .

3.2.1.2.1 Duration specification

In this section I discuss the specification of the conditional expectation ψ_i . As mentioned in the previous section, the range of potential ACD specifications is so wide that a complete review of those specifications is beyond the scope of this Chapter. The book of Hautsch (2004) provides a very exhaustive survey of existing ACD specifications. In Chapter 3 I focus on the standard ACD specification (Engle and Russell 1998) and the log-ACD of Bauwens and Giot (2000).

The conditional expectation of the duration ψ_i in a standard ACD (m, q) model is specified as follows,

$$\psi_i = \omega + \sum_{j=1}^m \alpha_j w_{i-j} + \sum_{j=1}^q \beta_j \psi_{i-j}. \quad (3.13)$$

(3.13) is a linear parameterization of (3.8). Specifically, the conditional expectation depends linearly upon the past durations and the past expected durations. In order to ensure that the conditional expectations to be nonnegative, parameters ω , α and β are required to have, $\omega > 0$, $\alpha \geq 0$ and $\beta \geq 0$. Obviously, the ACD model can be viewed as an analog of the GARCH (Bollerslev 1986) of the durations. The weak (covariance) stationarity conditions for the ACD model are just similar to those of the GARCH model. Specifically, $\sum_{j=1}^m \alpha_j + \sum_{j=1}^q \beta_j < 1$. A detailed discussion regarding the stationarity of ACD model can be found in Engle and Russell (1995). Under the assumption of weak stationarity, the unconditional mean and the conditional variance can be deduced from (3.9) and (3.13) straightforwardly,

$$E(w_i) = E(\psi_i)E(\epsilon_i) = \frac{\omega}{1 - \sum_{j=1}^m \alpha_j - \sum_{j=1}^q \beta_j} \quad (3.14)$$

$$Var(w_i | \bar{w}_{i-1}) = \psi_i^2 Var(\epsilon_i). \quad (3.15)$$

The unconditional variance of durations $Var(w_i^2)$ can also be deduced from (3.9) and (3.13). However, the general expression needs a direct computation of $E(\psi_i^2)$ that is rather cumbersome for a general ACD (m, q) . For the ACD $(1, 1)$ model, the unconditional variance of durations can be derived as follows,

$$E(w_i^2) = E(\psi_i^2)E(\epsilon_i^2) = \frac{\omega^2 E(\epsilon_i^2)}{1 - \beta^2 - 2\alpha\beta - \alpha^2 E(\epsilon_i^2)} + \frac{(\alpha + \beta)2\omega^2 E(\epsilon_i^2)}{(1 - \beta^2 - 2\alpha\beta - \alpha^2 E(\epsilon_i^2))(1 - \alpha - \beta)}, \quad (3.16)$$

$$E(w_i^2) - E^2(w_i) = E^2(w_i) \left[\frac{E(\epsilon_i^2)(1 - \beta^2 - 2\alpha\beta - \alpha^2) - (1 - \beta^2 - 2\alpha\beta - \alpha^2 E(\epsilon_i^2))}{1 - \beta^2 - 2\alpha\beta - \alpha^2 E(\epsilon_i^2)} \right]. \quad (3.17)$$

If ϵ_i is assumed to follow the exponential distribution, then $E(\epsilon_i^2) = 2$. Consequently, the unconditional variance in (3.17) can be simplified as

$$E(w_i^2) - E^2(w_i) = Var(w_i) = E^2(w_i) \left(\frac{1 - \beta^2 - 2\alpha\beta}{1 - \beta^2 - 2\alpha\beta - 2\alpha^2} \right). \quad (3.18)$$

(3.18) has an important implication that if $\alpha > 0$ the unconditional standard deviation of durations will be larger than the unconditional mean of durations. This phenomenon is usually referred to as “excessive dispersion”. For discussion and analysis regarding the

existence of higher moments of the ACD (m, q) model, see e.g., Carrasco and Chen (2002).

A useful feature of the ACD model is that it can also be formulated as the familiar form of an ARMA $(\max(m, q), q)$ model. Specifically, denote the martingale difference η_i as $\eta_i = w_i - \psi_i$, then the ACD (m, q) model can be written as

$$w_i = \omega + \sum_{j=1}^{\max(m, q)} (\alpha_j + \beta_j) w_{i-j} - \sum_{j=1}^q \beta_j \eta_j + \eta_i. \quad (3.20)$$

The ARMA-presentation in (3.20) can be used to derive the conditions for the stationarity and invertibility of the ACD model. The derivations and results are almost the same as work for GARCH models (Nelson 1992). The stationarity and invertibility requires that all roots of $1 - \alpha(L) - \beta(L)$ and $1 - \beta(L)$ lie outside the unit circle respectively, where $\alpha(L)$ and $\beta(L)$ are the polynomials in terms of the lag operator L .

The ARMA-presentation of the ACD model in (3.20) also can be used to derive the autocorrelation function of ACD models. For instance, the first order autocorrelation is now

$$\rho_1 = \text{Cov}(w_i, w_{i-1}) = \frac{\alpha_1(1 - \beta_1^2 - \alpha_1\beta_1)}{1 - \beta_1^2 - 2\alpha_1\beta_1}. \quad (3.21)$$

For more general description of the autocorrelation structure of the ACD model such as the Yule-Walker equations, Bauwens and Giot (2000) provided a detailed discussion.

As discussed previously, there are many other possible specifications for the conditional expectation of the duration ψ_i . Another important specification of the durations is the logarithmic-ACD model of Bauwens and Giot (2000). They argued that the imposed nonnegativity of parameters in (3.13) are too restrictive. For instance, if one wants to incorporate other economic explanatory variables into (3.13) with coefficients expected to be negative, the resulting conditional expectation of the duration might be negative. In a log-ACD (m, q) model, the autoregressive equation is specified using the logarithm of the conditional expectations of durations. Specifically,

$$\begin{aligned} w_i &= e^{\psi_i \epsilon_i}, \\ \ln \psi_i &= \omega + \sum_{j=1}^m \alpha_j \ln w_{i-j} + \sum_{j=1}^q \beta_j \ln \psi_{i-j}. \end{aligned} \quad (3.22)$$

Another logarithmic form of the conditional expectation of durations ψ_i is

$$\ln \psi_i = \omega + \sum_{j=1}^m \alpha_j \frac{w_{i-j}}{\psi_{i-j}} + \sum_{j=1}^q \beta_j \ln \psi_{i-j}. \quad (3.23)$$

The logarithmic form of the conditional expectations of durations ψ_i in (3.22) and (3.23) releases the nonnegative constraints of parameters α and β in the standard ACD model. In order to ensure the weak stationarity, (3.22) requires that $|\alpha + \beta| < 1$ while (3.23) requires that $|\beta| < 1$. Empirical evidence in general favors the specification of (3.23). See, e.g., Bauwens and Giot (2000) and Dufour and Engle (2000).

It is possible to derive the unconditional moments and the autocorrelation function of the log-ACD model from (3.22) and (3.23). However, the expression is rather cumbersome and therefore will not be presented here. Bauwens, Galli, and Giot (2003) presented a detailed derivation. Compared with the standard ACD model that has an autocorrelation function decreasing geometrically at the rate $\alpha + \beta$ (see (3.20)), the log-ACD model has an autocorrelation function that decreases at a rate less than β for small lags (Bauwens, Galli and Giot, 2003). As for the specification of ϵ_i , Bauwens and Giot (2000) assumed that it follows the Weibull distribution. In general, many other candidates of the density $p(\epsilon; \phi)$ can be chosen for ϵ_i as in the standard ACD model. The maximum likelihood estimation of the log-ACD is an analog of the maximum likelihood estimation of the standard ACD models.

I will elaborate the possible candidates of the distributional assumption of ϵ_i and present the maximum likelihood estimation of the ACD model in the next section.

3.2.1.2.2 Distributional assumption and maximum likelihood estimation

As discussed in the previous section, an important feature of the ACD model is that the specification in (3.9) allows for flexibility in the selection of the standardized duration ϵ_i . In general, any distribution with nonnegative supports could be used in the modelling of durations. Among the vast possible candidates, I focus on the following three distributions that are commonly applied in practice: exponential distribution, Weibull distribution and generalized gamma distribution.

It seems natural to start with the exponential distribution since it is widely used in the modelling for the “waiting-time” in many fields. Although the assumption of exponentially distributed standardized durations seems restrictive in many circumstances, it still serves as a good starting point since it can provide quasi-maximum likelihood estimators for the ACD parameters.

Under the assumption that the standardized durations are exponentially distributed, the log-likelihood function can be derived from (3.11) and can be presented as

$$L(\theta_\psi, \phi) = -\sum_{i=1}^{N(T)} \left[\frac{w_i}{\psi_i} + \ln \psi_i \right]. \quad (3.24)$$

I use the terminology “Exponential-ACD” (EACD) to refer to the standard ACD model with exponentially distributed standardized durations ϵ_i . The strong analogy between the EACD and the GARCH model can be further verified by the similar quasi-maximum likelihood estimation (QMLE) properties of the two models. Specifically, the QMLE properties of the GARCH (1,1) model derived by Lee and Hansen (1994) and Lumsdaine (1996) also holds for the EACD (1,1) model. Engle and Russell (1998) and Haustsch (2004) provided detailed proof for the equivalence of QMLE property of the two models. Specifically, Theorem 1 and Theorem 3 in Lee and Hansen (1994) can be rewritten in the ACD form as:

Assume the following conditions:

- (i) Θ is a compact parameter space and $\text{int } \Theta \neq \emptyset$.
- (ii) $\theta_0 = (\omega_0, \alpha_0, \beta_0) \in \text{int } \Theta$.
- (iii) $E(w_i | F_{i-1}; \theta_0) \equiv \psi_i(\theta_0) \equiv \psi_{i,0} = \omega_0 + \alpha_0 w_{i-1} + \beta_0 \psi_{i-1}$.
- (iv) $\epsilon_i = \frac{w_i}{\psi_i}$ is strictly stationary, ergodic and nondegenerate.
- (v) $E(\epsilon_i^2 | F_{i-1}) < \infty$ almost surely.
- (vi) $\sup_i E[\ln \beta_0 + \alpha_0 \epsilon_i | F_{i-1}] < 0$ almost surely.
- (vii) $L(\theta) = \sum_{i=1}^n l_i(\theta) = -\sum_{i=1}^N \left[\frac{w_i}{\psi_i} + \ln \psi_i \right]$, where $\psi_i = \omega + \alpha w_{i-1} + \beta \psi_{i-1}$.

Then the maximizer of L is consistent and asymptotically normal. The associated covariance matrix is given by the robust standard errors by Lee and Hansen (1994).

The above corollary of the Theorem 1 and Theorem 3 in Lee and Hansen (1994) has several important implications. First, it illustrates that the maximization of the log likelihood function in (3.24) yields a consistent estimate of θ without a clear specification of the standardized duration ϵ_i . In other words, the estimates of parameters of the EACD (1,1) is consistent and asymptotically normal even under misspecification of ϵ_i . Consequently, the EACD (1,1) can be directly estimated by GARCH software. The procedure is straightforward: use $\sqrt{w_i}$ as the dependent variable in the GARCH regression equation and set the conditional mean as zero.

Second, when $\alpha + \beta < 1$, the condition (vi) is satisfied, but this inequality is a sufficient but not necessary condition for the Theorem to hold. Thus, it may hold for integrated duration process in which $\alpha + \beta = 1$. Third, although it requires the strict stationarity and ergodicity of ϵ_i , it does not require ϵ_i to follow an iid process. Hence, it establishes the QMLE properties for many other ACD specifications such as the semiparametric ACD model of Drost and Werker (2001).

Nevertheless, it is noticeable that the results are based on the EACD (1,1) model and cannot be extended necessarily to the general EACD (m, q) model. Besides, as suggested by the condition (iii), the QMLE property of the EACD (1,1) model depends largely on the correct specification of the conditional expectation ψ_i . Moreover, empirical research in general are against the assumption of exponentially distributed standardized duration ϵ_i (Gramming and Maurer 2000, Feng et al. 2004, Lin and Tamvakis 2004 and Dufour and Engle 2000). Thus, if the data set is relatively small, the quasi-maximum likelihood estimates of the parameters of the EACD (1,1) model might be seriously biased. Finally, the hazard function of the exponential distribution is a constant and therefore is clearly independent of the time passed since the last trade. This characteristic is undesirable in the modelling of ultra-high-frequency transactions for obvious reasons. In order to resolve these issues, it seems natural to choose a more general distribution of the standardized durations ϵ_i and to replace the quasi-maximum likelihood estimation by standard maximum likelihood estimation. This procedure is similar to the case of ARCH-GARCH model where the normal assumption of returns is often replaced by some leptokurtic distributions.

A popular alternative of the exponential distribution is the Weibull distribution. (Engle and Russell 1998 and Bauwens and Giot 2000). Consider a Weibull density $p(\epsilon; \lambda, k)$ for the standardized duration ϵ_i where λ is the scale parameter and k is the shape parameter. This assumption is equivalent to assuming that ϵ_i^k is exponentially distributed. I use the terminology “Weibull-ACD” (WACD) to refer to this class of ACD models. The Weibull distribution reduces to the exponential distribution when $k = 1$. The probability density of the Weibull distribution can be presented as

$$p(\epsilon; \lambda, k) = \frac{k}{\lambda} \left(\frac{\epsilon}{\lambda}\right)^{k-1} \exp\left(-\frac{\epsilon}{\lambda}\right)^k, \quad (3.25)$$

where $\epsilon > 0, \lambda > 0$ and $k > 0$. The hazard function of the Weibull distribution can be presented as

$$h(\epsilon) = \left(\frac{1}{\lambda}\right)^k \epsilon^{k-1} k. \quad (3.26)$$

Under the ACD framework, the conditional intensity of Equation (3.12) can be presented as

$$\lambda(\tau | N(\tau), t_0, \dots, t_{N(\tau)}) = \left[\frac{\Gamma(1 + \frac{1}{k})}{\psi_{N(\tau)+1}} \right]^k (\tau - t_{N(\tau)})^{k-1} k, \quad (3.27)$$

where $\Gamma(\cdot)$ is the gamma function and k is the shape parameter of the Weibull distribution. Observe that the conditional intensity function in (3.27) is monotonic with respect to τ as the hazard function in (3.26). It is increasing if $k > 1$ and is decreasing if $k < 1$. As a result, this monotonic property of the Weibull hazard function makes the occurrence of long durations, compared with exponential distributions, more or less likely dependent on whether $k < 1$ or $k > 1$. The log-likelihood function under the Weibull distribution can be derived from (3.11) and (3.13). Specifically, assume the scale parameter λ to be unit for the simplicity of presentation, then

$$L(\theta_\psi, k) = \sum_{i=1}^{N(T)} \left[\ln\left(\frac{k}{\psi_i}\right) + k \ln\left(\frac{\Gamma(1 + \frac{1}{k}) \psi_i}{\psi_i}\right) - \left(\frac{\Gamma(1 + \frac{1}{k}) \psi_i}{\psi_i}\right)^k - \ln \Gamma\left(1 + \frac{1}{k}\right) \right]. \quad (3.28)$$

When $k = 1$, (3.28) reduces to the (3.24).

The last family of distributions that I consider in this chapter is the generalized gamma distribution. The generalized gamma distribution nests the Weibull distribution, the Exponential distribution and the Gamma distribution and therefore provides greater

flexibility in the modelling of standardized duration ϵ_i . In 1996, Lunde introduced the use of generalized gamma distribution in the modelling of durations. The density of the generalized gamma distribution is

$$p(\epsilon; a, d, m) = \frac{m\epsilon^{d-1}}{a^d \Gamma(d/m)} \exp\left(-\frac{\epsilon}{a}\right)^m. \quad (3.29)$$

where $\Gamma(\cdot)$ is the gamma function and $a > 0$, $d > 0$, $m > 0$ and $\epsilon > 0$.

One can immediately infer from the density of generalized gamma distribution in (3.29) that when $d = m$ the density reduces to the density of the Weibull distribution. Further assume that $d = m = 1$, the density now is the density of the exponential distribution. Besides, if $m = 1$, the density reduces to the density of the Gamma distribution. Glaser (1980) provided a discussion of the relations between the parameters and the shape of hazard function for the generalized gamma distribution. The concrete presentation of the hazard function of the generalized gamma distribution is complicated since it involves an incomplete gamma integral and therefore has no closed form. Numerical techniques are usually applied to depict the pattern of the hazard functions. For further details regarding the hazard function of the generalized gamma distribution, see e.g., Glaser (1980). An important property of the hazard function of the generalized gamma distribution is that it is not monotonic. For instance, given that $dm < 1$, if $m > 1$ the hazard function has a U-shape and else if $m \leq 1$ the hazard function is decreasing. Given that $dm > 1$, if $m > 1$ the hazard function is increasing and else if $m \leq 1$ the hazard function has an inverted U-shape. This property of the hazard function of generalized gamma distributions allows for more flexibility in the modelling of the instantaneous rate of occurrence of transactions.

The derivation of the log-likelihood function of “Generalized Gamma-ACD” (GG-ACD) from (3.11) is straightforward, substituting the density function in (3.11) with (3.29),

$$L(\theta_\psi, a, d, m) = \sum_{i=1}^{N(T)} \left[\ln \frac{m}{aw_i} + d \ln \frac{w_i}{\psi_i} \frac{\Gamma(\frac{d+1}{m})}{\Gamma(\frac{d}{m})} + \left(-\frac{w_i}{\psi_i} \frac{\Gamma(\frac{d+1}{m})}{\Gamma(\frac{d}{m})} \right)^m - \ln \Gamma\left(\frac{d+1}{m}\right) \right]. \quad (3.30)$$

The detailed derivation could be found in Lunde (1996). Observe that (3.30) reduces to (3.29) if $a = 1$ and $d = m$. It further reduces to equation (3.24) if $a = m = d = 1$.

To conclude this section, I present the conditional density of w_i dependent on past information \bar{w}_{i-1} with respect to the generalized gamma distribution under the general ACD framework of (3.7) - (3.9) for the completeness of the introduction.

$$f(w_i | \bar{w}_{i-1}; a, d, m, \theta_\psi) = \frac{m}{a w_i \Gamma(\frac{d+1}{m})} \left[\frac{w_i}{\psi_i} \frac{\Gamma(\frac{d+1}{m})}{\Gamma(\frac{d}{m})} \right]^d \exp \left(- \left[\frac{w_i}{\psi_i} \frac{\Gamma(\frac{d+1}{m})}{\Gamma(\frac{d}{m})} \right] \right)^m. \quad (3.31)$$

3.2.1.3 Durations and marks

It is often the case that we are not only interested in the durations between transactions but also the marks associated with the transactions. The ACD models discussed previously are proposed for describing the dynamics of durations given past information. As suggested by (3.5), they describe the marginal density of the duration w_i . Thus, it remains to present a framework that could jointly describe the durations and marks. Especially, it is often the case that given the durations and the past information, one wants to test some market microstructure hypothesis.

In the Section 3.2.1.3.1, I present and discuss a framework of modelling tick-by-tick transaction dynamics based on the point process. Section 3.2.1.3.2 focus on the volatility modelling with in the ACD framework. Section 3.2.1.3.3 discusses the incorporation of possible explanatory economic variables.

3.2.1.3.1 General setup

As discussed previously, I view the arrival of transactions as a point process. Specifically, denote the timing of transactions as $\{t_0, t_1, \dots, t_i, \dots, t_N\}$ and denote the i th duration between two transactions that occur at t_{i-1} and t_i as $w_i = t_i - t_{i-1}$. Let z_i be the marks such as volume and prices associated with the i th transaction. The durations and marks can be jointly presented as the sequence $\{(w_i, z_i)\}_{i=1,2,\dots,N}$. Thus, the joint conditional density of the i th element (w_i, z_i) can be presented as

$$(w_i, z_i) | F_{i-1} \sim f(w_i, z_i | w_1, w_2, \dots, w_{i-1}, z_1, z_2, \dots, z_{i-1}; \theta_i). \quad (3.32)$$

(3.32) is a slight modification of the (3.4). The main difference is that the parameters θ_i 's are now conditional on the past information and therefore can be potentially different for different transaction observations. The conditional density can be rewritten as the product of marginal conditional density of the durations and the conditional density of the marks as in (5).

$$f(w_i, z_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_i) = f_w(w_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_{w,i}) g(z_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_{z,i}), \quad (3.33)$$

where $\bar{w}_{i-1} = (w_1, w_2, \dots, w_{i-1})$ and $\bar{z}_{i-1} = (z_1, z_2, \dots, z_{i-1})$ denote the past information of w_i and z_i . In the real world, the realized transactions are nearly zero at every point in the time. Therefore, it is natural to measure the probability of the occurrence of transactions at specific time in terms of the limit. The conditional intensity function of (3.1) are usually applied to describe such probabilities. Under the specification of (3.32) and (3.33), then conditional intensity of (3.1) can be rewritten as

$$\lambda_i(t | \bar{w}_{i-1}, \bar{z}_{i-1}) = \lim_{h \rightarrow 0^+} \frac{\Pr(N(t+h) > N(t) | \bar{w}_{i-1}, \bar{z}_{i-1})}{h}, t \in [t_{i-1}, t_i]. \quad (3.34)$$

$N(t)$ is the function of the number of transactions that occur by time t . Another presentation of (3.34) can be obtained by using (3.33). For $t \in [t_{i-1}, t_i]$, the probability of a transaction should be conditional to both past information such as \bar{w}_{i-1} and \bar{z}_{i-1} and the information that there has been no transaction since t_{i-1} . This is usually expressed as the ratio of probability density function and survival function. By (3.33) and (3.34), one can have

$$\lambda_i(t | \bar{w}_{i-1}, \bar{z}_{i-1}) = \frac{f_w(t - t_{i-1} | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_{w,i})}{\int_{\tau \geq t} f_w(\tau - t_{i-1} | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_{w,i}) d\tau}, t \in [t_{i-1}, t_i]. \quad (3.35)$$

Many questions of economic interests can be explored in the setup of (3.32) - (3.35). For instance, the conditional marginal density of the next mark z_i given \bar{z}_{i-1} is

$$f_z(z_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_i) = \int_{\tau} f(w_i, z_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_i) d\tau. \quad (3.36)$$

Once the f_w and g in (3.33) are specified, the log-likelihood function can be obtained immediately,

$$L(\theta) = \sum_{i=1}^N \ln f(w_i, z_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_i)$$

$$= \sum_{i=1}^N [\ln f_w(w_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_{w,i}) + \ln g(z_i | \bar{w}_i, \bar{z}_{i-1}; \theta_{z,i})]. \quad (3.37)$$

The log-likelihood function in (3.37) should be maximized jointly by the sum of the first term of $\ln f_w(w_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_{w,i})$ and the second term of $\ln g(z_i | \bar{w}_i, \bar{z}_{i-1}; \theta_{z,i})$. Separating the maximization of the first term and the second term comes at the cost of the efficiency of the MLE. In practice, a two-step procedure is usually adopted for a large sample. Specifically, one first takes the estimates of θ_w by maximizing the first term in (3.37) and further gives the θ_w into the maximization of second term in (3.37) to obtain the estimates of θ_z . The estimates in the two-step estimation are consistent but inefficient. The main reasoning behind this procedure is that, given a large sample, the loss of efficiency might be acceptable (Engle 2000).

3.2.1.3.2 Durations and Volatility

Among the marks contained in z_i , the most important information is the execution price p_i associated with the transaction. The price reflects the volatility of the underlying asset and the market. In this section we first discuss the relations between durations and volatility. I review the market microstructure theory that outlines the possible explanation of the relation between durations and volatility such as Diamond and Verrecchia (1987), Easley and O'Hara (1992), and Admati and Pfleiderer (1988).

Further, I analyze the GARCH volatility modelling under the ACD framework. The Ultra-High-Frequency (UHF) GARCH volatility model of Engle (2000) will be presented as an example. The joint modelling of durations and marks can also be obtained by exploiting the conditional intensity function of (3.35). A detailed introduction can be found in Hautsch (2004).

In 1987, Diamond and Verrecchia argued that “no trade means bad news”. The theory can be briefly explained by the following scenario. Suppose there are some informed traders who currently do not own any stock and are prevented from short selling. Thus, they cannot profit from the asymmetric information and will refuse to trade at the current market prices. The specialist whose role is to facilitate trading for certain stocks for some stock exchanges such as NYSE learns from the durations and further lower their prices. Consequently, the

existence of no transactions can be viewed as a signal of bad news.

Easley and O'Hara (1992) provided another economic interpretation of the durations between transactions. They focused on the arrival of new information. Suppose there are some informed traders who know whether there is new information or not. They will buy or sell depending on whether the information is good or bad respectively. Thus, long durations between transactions can be interpreted as the evidence of no news. Their theory therefore can be briefly described as “no trade means no news”. Thus, under the framework of their theory the durations should be negatively correlated with volatilities.

Admati and Pfleiderer (1988) gave a very interesting alternative point of view. They argued that the short durations between transactions are the result of some liquidity traders. The trading behavior of these liquidity traders is not based on asymmetric information. As a result, the volatility should be low when durations are short. In contrast, relatively long durations could be interpreted as a sign of the leaving of these liquidity traders. Consequently, there will be a high proportion of informed traders. The volatility should be high since these transactions reflect the arrival of new information. Their theory suggests that the durations should be positively correlated with volatilities.

It is worthy to mention that the literatures of Easley and O'Hara (1992) and Admati and Pfleiderer (1988) are more relevant to the liquidity rather than volatility. The discussion above is an attempt to connect their microstructure theories with the ACD modelling of durations and volatility.

Given the theoretical background provided by above market microstructure literatures, it remains to present a parametric framework to jointly describe prices volatility and durations. For a tick-by-tick transaction data set, let the sequence of execution prices associated with each transaction be $\{p_0, p_1, \dots, p_i, \dots, p_N\}$. The corresponding transaction time is $\{t_0, t_1, \dots, t_i, \dots, t_N\}$. The i th return r_i and the i th duration w_i are defined as $r_i = \ln \frac{p_i}{p_{i-1}}$ and $w_i = t_i - t_{i-1}$. Thus, the tick-by-tick return series and its corresponding duration series are $\{r_1, \dots, r_i, \dots, r_N\}$ and $\{w_1, \dots, w_i, \dots, w_N\}$ respectively.

With respect to the volatility modelling, since the return series are defined as $\{r_1, \dots, r_i, \dots, r_N\}$, one can therefore define the conditional variance for return r_i as

$$v_{i-1}(r_i | \bar{w}_i, \bar{r}_{i-1}) = q_i. \quad (3.38)$$

The definition of the conditional variance q_i in (3.38) suggests that the conditional variance for the i th return depends on not only the past information of returns \bar{r}_{i-1} but also on the current and past information of durations \bar{w}_i . Besides, (3.38) defines the conditional variance associated with return r_i from t_{i-1} to t_i . This definition is associated with each transaction rather than fixed time interval. Specifically, note that q_i and q_j now represent the conditional variance measured from durations with different length if $i \neq j$. In practice, it is often the case that the tick-by-tick durations are multiples of the minimum unit of time measurement. Consequently, the durations present a salient feature of discreteness. It seems convenient to transform the conditional variance for each return into the conditional volatility per unit of time. By (3.38), the transformation can be presented as

$$v_{i-1}\left(\frac{r_i}{\sqrt{w_i}} \middle| \bar{w}_i, \bar{r}_{i-1}\right) = \sigma_i^2. \quad (3.39)$$

Under the definition of (3.38) and (3.39), the relation between q_i and σ_i can be presented as

$$q_i = w_i \sigma_i^2. \quad (3.40)$$

Thus, given past information of \bar{w}_{i-1} and \bar{r}_{i-1} , the predicted expectation of q_i is

$$E_{i-1}(q_i | \bar{w}_{i-1}, \bar{r}_{i-1}) = E_{i-1}(w_i \sigma_i^2 | \bar{w}_{i-1}, \bar{r}_{i-1}). \quad (3.41)$$

(3.40) and (3.41) suggest a simple methodology of incorporating durations into standard techniques of volatility modelling. Since σ_i^2 now is the conditional variance per unit time, standard time series techniques based on fixed time interval such as GARCH volatility model can be applied to describe its dynamics. Specifically,

$$\begin{aligned} \frac{r_i}{\sqrt{w_i}} &= \sigma_i z_i, \\ \sigma_i^2 &= \omega + \alpha \frac{r_{i-1}^2}{w_{i-1}} + \beta \sigma_{i-1}^2, \end{aligned} \quad (3.42)$$

where z_i is an iid sequence with zero mean and unit variance. (3.42) is the simplest form of the UHF-GARCH of Engle (2000). It can be viewed as a modification of the standard GARCH (1,1) model that is adjusted to account for the irregularly spaced transactions. Recall the ACD specification of w_i in (3.9),

$$w_i = \psi_i \epsilon_i,$$

where ψ_i is the conditional expectation of w_i dependent on past information of durations \bar{w}_{i-1} . A substitution of (3.9) into the (3.42) gives

$$r_i = \sigma_i z_i \sqrt{\psi_i \epsilon_i}. \quad (3.43)$$

Observe that under (3.42), the current duration is not informative since σ_i^2 depends on only past information of durations and returns. The implicit assumption behind this specification is that the news carried by the last return r_{i-1} is fully captured by the square of the last return innovations σ_{i-1} per unit of time and is independent of current durations. However, as suggested by the market microstructure papers, the variation of durations and variation of return volatilities might be related. In order to incorporate the information of current duration, the volatility specification of (3.42) can be modified to account for current durations. For instance, Engle (2000) proposed the following specification of σ_i^2 .

$$\sigma_i^2 = \omega + \alpha \frac{r_{i-1}^2}{w_{i-1}} + \beta \sigma_{i-1}^2 + \gamma \frac{1}{w_i}. \quad (3.44)$$

(3.44) corresponds to the theory of Easley and O'Hara (1992). The reciprocal of the durations in (3.44) suggests that long durations contribute less to the variance. This implication is in line with the argument of “no trade means no news”. Observe that under (3.44), the current return r_i now depends on current durations and past durations and prices as in (3.33). The joint conditional density $f(w_i, z_i | \bar{w}_{i-1}, \bar{z}_{i-1}; \theta_i)$ in (3.33) therefore can be determined.

Another model that describes the price dynamics based on ACD modelling of durations is the Autoregressive Conditional Multinomial (ACM) model of Russell and Engle (2005). Building on (3.5), they applied an ACD specification for durations and a dynamic multinomial model for distribution of price changes conditional on past information and the current duration. Nevertheless, the original ACM model (Engle and Russell 2005) depends on a relatively large number of states that characterize different regimes of durations dynamics for different price changes. Consequently, many parameters are needed, which inevitably complicates the estimation. For a two-state ACM model, see Prigent et al. (2001).

Finally, other economic explanatory variables can be introduced into the conditional volatility specification in (3.42) such as the bid-ask spread and volumes to provide a more realistic modelling of the conditional volatility. For instance, Engle (2000) included the trading intensity and volume in (3.42) and found that compared with the trading intensity, volume has very little explanatory power over the conditional variance σ_i^2 .

3.2.1.3.3 Durations and other explanatory variables

In this section I discuss the relation between durations and other microstructure variables. I first present the incorporation of explanatory variables into the ACD model. Then, we discuss the market microstructure variables of interests by reviewing the market microstructure literatures with respect to the timing of transactions.

Due to the flexibility of the ACD specification, economic explanatory variables also can be introduced into the modelling of durations. In general, there are two different ways to incorporate the possible explanatory variables.

First, one can add the random variable of interests directly into the specification of conditional expectation ψ_i . Specifically, recall (3.13) that gives the conditional expectation of the duration ψ_i for a standard ACD (m, q) model,

$$\psi_i = \omega + \sum_{j=1}^m \alpha_j w_{i-j} + \sum_{j=1}^q \beta_j \psi_{i-j}.$$

Suppose the explanatory variables that we are interested in is ξ_i . Include this variable in (3.13), then

$$\psi_i = \omega + \sum_{j=1}^m \alpha_j w_{i-j} + \sum_{j=1}^q \beta_j \psi_{i-j} + \gamma \xi_{i-1}. \quad (3.45)$$

An alternative form of incorporating ξ_i into (3.13) is,

$$\psi_i - \gamma \xi_{i-1} = \omega + \sum_{j=1}^m \alpha_j w_{i-j} + \sum_{j=1}^q \beta_j (\psi_{i-j} - \gamma \xi_{i-1} - j). \quad (3.46)$$

As discussed previously, (3.45) and (3.46) cannot necessarily guarantee the nonnegativity of the conditional expectation of durations ψ_i if ξ_i can take negative values. There are two common solutions to this issue. First, one can consider presenting ξ_i in some function forms that are nonnegative in nature such as the exponential function. Second, one can use the Log-ACD of Bauwens and Giot (2000). The Log-ACD (m, q) specification is given by

(3.22),

$$w_i = e^{\psi_i \epsilon_i},$$

$$\ln \psi_i = \omega + \sum_{j=1}^m \alpha_j \ln w_{i-j} + \sum_{j=1}^q \beta_j \ln \psi_{i-j}.$$

or

$$\ln \psi_i = \omega + \sum_{j=1}^m \alpha_j \frac{w_{i-j}}{\psi_{i-j}} + \sum_{j=1}^q \beta_j \ln \psi_{i-j}.$$

Observe that the use of logarithmic form of the conditional expectations of durations ψ_i releases the nonnegative constraints of parameters α and β in the standard ACD model. Similarly, incorporating explanatory variables into the Log-ACD model now will not suffer from the possibly negative conditional expectations of durations.

The second way of including explanatory variables is to consider the explanatory variable as a scaling function of the durations w_i . For instance, denote the explanatory variable of interest as ξ_i . Let h be some nonnegative function. Then

$$w_i = h(\xi_{i-1}) \psi_i \epsilon_i. \quad (3.47)$$

Observe that the variable ξ in (3.47) is indexed by $i-1$. It is due to the definition that $w_i = t_i - t_{i-1}$. Thus, w_i are subject to information until t_{i-1} . Thus, the index is labeled with one lag in order to corresponds to sequence of transaction timing $\{t_1, t_2, \dots, t_i, \dots\}$. One can of course use ξ_i if the indexes are given corresponding to the sequence of durations $\{w_1, w_2, \dots, w_i, \dots\}$.

A simple and illustrative example is the inclusion of intraday periodicity. Denote the intraday periodicity of the durations as s_i . Then it can be incorporated into the ACD modelling of durations by

$$w_i = s_{i-1} \psi_i \epsilon_i. \quad (3.48)$$

I now discuss the possible economic variables that can be introduced into the ACD modelling. A natural candidate is the volume associated with each transaction, the relationship between volume and price dynamics is widely identified by both theoretical and empirical research such as Easley and O'Hara (1992), Easley et al. (1997) and Blume et al. (1994). These papers suggest that, in general, volumes convey information which are not fully captured by price dynamics. For instance, according to Easley and O'Hara (1992),

the abnormally large volume might be viewed as the signal of the arrival of informed traders. In general, empirical evidence (Engle 2000, Engle and Russell 1998 and Dufour and Engle 2000) found that the volume is negatively correlated with the conditional expectation of duration ψ_i .

Another market microstructure variable of particular interests is the bid-ask spread. Easley and O'Hara (1992) argued that a high spread might be an indication of informed trading. Consequently, it should be related with the short durations. The reasoning is that under the framework of Easley and O'Hara (1992), the volatility of prices will increase if the proportion of informed traders increases. Consequently, the bid-ask spread will be widened. Thus, the wide bid-ask spread could be informative of information arrivals. For a theory of "no trade means no news", the wide spread therefore indicates the new information and should be related with short durations. In order to analyze this effect, the spread for each transaction can be considered as a random variable and can be incorporated into the duration modelling. The empirical evidence of Bauwens and Giot (2000) and Engle (2000) suggest that the bid-ask spread is negatively correlated to the durations.

As discussed in Chapter 1, compared with volumes, the number of transactions that occur during a specific time interval has significantly higher explanatory powerful over the price volatility during the time interval (Ane and Geman 2000, Easley et al. 1995 and Jones et al. 1994). Jone et al. (1994) provided a market microstructure theory for this phenomenon. They argued that the number of transactions is closely related to the change of prices since transactions are more likely to happen when there is new information. In contrast, the framework of Admati and Pfleiderer (1988) suggests that the intensity of transaction should be negatively correlated with volatility.

The duration itself, of course, is a measure of trading intensity since it measures how frequently an equity is traded. For a tick-by-tick transaction record, the corresponding duration series directly reflect the transaction intensity. To analyze the effect of the trading intensity, it is necessary to conduct a deformation of the original durations. Besides, it has been reported that a large percentage of the tick-by-tick transactions has no price changes

(Tsay 2002, Bertram 2005 and Dionne et al., 2005).

Intraday transaction prices often can take only finite values due to some institutional features regarding the price restrictions. For instance, financial markets like NYSE usually specifies a minimum unit of price measurement, called a tick. In other words, transaction prices must fall on a grid. Given such a grid, the durations (Engle and Russell 1998 and Giot 1999) can be defined by filtering the tick-by-tick transactions and retaining those leading to a significant change of prices. The durations defined in this sense are called the price durations. Observe that the price durations are of different length and may contain different number of tick-by-tick transactions. Thus, for example, a large number of tick-by-tick transactions over a short duration represents a high trading intensity. Clearly, volume durations can be defined in a similar way. For a price duration, trading intensity can be defined as ratio of the number of transactions contained in the duration to the length of the durations. (3.45) and (3.46) now can be applied to analyze the effect of the trading intensity. Giot and Bauwens (2000) argued that there is a highly significant negative correlation between trading intensity and the expectation of conditional durations.

Given the vast existing market microstructure literatures with respect to the relation between durations and associated market microstructure variables such as volume and prices, it is widely accepted that time is not exogenous to the price process. Nevertheless, the interdependence of those variables is still open to question. Under the framework of marked point process, such interdependence can be studied through the decomposition of (3.33). For instance, one can add volumes and bid-ask spread to the conditional volatility specification of the UHF-GARCH model to examine their explanatory power over price dynamics. A wide range of closely related model following this framework can be found in Gramming and Wellner (2002), Ghysels and Jasiak (1998), Bauwens and Giot (2003) and Russell and Engle (2005).

An alternative approach for studying the interdependence between durations and other microstructure variables is provided by Hasbrouck (1991). He applied the vector autoregressive (VAR) system to analyzes the impact of current execution prices on the

future prices and presented a bivariate model for the relation between price changes and trade dynamics such as the sign of price movement. The timing of transactions is not considered as informative. Thus, the conditional density of the marks, $g(z_i | \bar{w}_i, \bar{z}_{i-1}; \theta_{z,i})$ in (3.33) depends on past information of z_i , which is $g(z_i | \bar{w}_i, \bar{z}_{i-1}; \theta_{z,i}) = g(z_i | \bar{z}_{i-1}; \theta_{z,i})$.

Dufour and Engle (2000) extended the model of Hasbrouck (1991) to account for the influence of durations on the price dynamics. They first applied an ACD model to describe the duration dynamics. Further, the duration is considered as a predetermined variable and therefore the coefficients for price changes and trade dynamics are allowed to be time varying. The model can be extended to incorporate marks of interest such as volume and spread easily. See e.g., Spierdijk (2004) and Manganelli (2005).

In order to capture the complicated interdependence between durations and marks, most of the joint models of durations and marks discussed above require a relatively large number of parameters. Thus, the maximization of the log-likelihood function of (3.37) faces the computation difficulties. In order to simplify the estimation procedures, it is usually assumed that durations have some form of exogeneity (Engle et al. 1983).

3.2.2 ACD in applications

In this section I present a discussion of the ACD models in application. As discussed previously, the ACD models are based on the point process modelling of transaction arrivals and can be applied to test many market microstructure hypotheses. It could be also used for the modelling of arrivals of a variety of financial events. In Section 3.2.2.1 I present the discussion of intraday periodicity of the intraday durations. As discussed in Chapter 2, like volatility and volume, the durations also present a seasonal pattern during the trading day. Section 3.2.2.2 include the test procedures of the ACD modellings. Section 3.2.2.3 reviews the literatures of the ACD models. I categorize the literatures according to the different definitions of durations.

3.2.2.1 Intraday Periodicity

As discussed in Chapter 2, the intraday transactions during the trading day are characterized by a strong periodicity. Early research using intraday data sampled from fixed time interval (the time-aggregated returns) focused on the intraday periodicity pattern in the price dynamics (Bollerslev and Andersen 1997,1998 and Beltratti and Morana 1998 and Engle 2000). In Chapter 2, I present an estimation of the intraday volatility periodicity that can be extended to returns that are irregularly spaced in time (the transaction-aggregated returns). The intraday volatility periodicity is found to have strong impact on the GARCH modelling of price dynamics.

With respect to the durations, Engle and Russell (1998) found that in general the transactions occurring at the beginning and closing hours of the trading day are associated with short durations. In contrast, the transactions that occur in the middle of the day are with relatively long durations. This finding might not be very surprising since it is well-known that compared with the middle day, the opening and closing trading hours are with much higher trading intensity. Consequently, the durations of transactions occurring in the opening and closing hours are expected to have shorter durations. These intraday periodic patterns of durations are often explained by the institutional features and trading habits. For instance, traders are highly active during the opening hours because they want to adjust their position according to the overnight news. Similarly, at the closing, traders want to close their position based on the information during the trading day. Lunchtime are related with less trading intensity due to its nature.

The findings of Engle and Russell (1998) provide direct evidence of the existence of the intraday duration periodicity. Thus, when modelling the durations dynamics under the ACD framework, the effect of such intraday duration periodicity should be included and analyzed since ignoring these intraday patterns might distort the estimation seriously. According to the results that I obtain in Chapter 2, the intraday volatility periodicity induces a very strong U-shape pattern into the autocorrelation of intraday volatilities. Standard GARCH volatility models cannot capture this effect and therefore often give estimations contradictory to the temporal aggregation predictions. Given the resemblance between the

ACD model and the GARCH model, it seems natural to exclude the intraday pattern of durations before applying the ACD model as in the case of GARCH intraday volatility modelling.

The fundamental technique to capture the effect of intraday pattern is very similar to the one that I use in the Chapter 2. This technique is proposed by Andersen and Bollerslev (1997,1998) for incorporating the intraday pattern into the intraday volatility modelling. Engle and Russell (1998) extended it to the duration modelling. Specifically, one can decompose the intraday durations as the product of a deterministic component that accounts for the intraday periodicity pattern and a stochastic component that models the duration dynamics. Let $s(t_i)$ be the intraday periodicity component at t_i and let $w_i = t_i - t_{i-1}$ be the i th duration. Then,

$$w_i = \tilde{w}_i s(t_{i-1}). \quad (3.49)$$

The \tilde{w}_i in (3.49) can be interpreted as the seasonal normalized durations. Now the conditional expectation ψ_i of w_i in (3.7) can be written as

$$\psi_i = E(w_i | \bar{w}_{i-1}, \bar{z}_{i-1}) s(t_{i-1}). \quad (3.50)$$

There are two approaches that are commonly used to specify the deterministic intraday periodicity component $s(t_i)$ (Andersen and Bollerslev 1997, Engle and Russell 1998, Bauwens and Giot 2000 and Engle 2000). The first one is to use a piecewise linear or cubic spline function to describe the intraday periodicity of durations. Further, the seasonal normalized durations \tilde{w}_i can be obtained by taking the ratio between the raw durations and the corresponding function values. For instance, one can first average the durations over ten minutes intervals for each trading day. Further, the mean of durations for each ten minutes interval can be estimated with the whole sample. Cubic splines are then applied on these intervals to smooth the intraday periodicity function.

The second approach involves the use of flexible Fourier series approximation. See, e.g., Andersen and Bollerslev (1998). This approach is based on the work of Gallant (1981). Specifically, the intraday periodic trend $s(t)$ is specified to follow the form of

$$s(t) = s(\delta^s, \bar{t}, q) = \bar{t}\delta^s + \sum_{j=1}^q [\delta_{c,j}^s \cos(\bar{t}2\pi j) + \delta_{s,j}^s \sin(\bar{t}2\pi j)], \quad (3.51)$$

where δ^s , $\delta_{c,j}^s$ and $\delta_{s,j}^s$ are the intraday seasonal coefficients that need to be estimated. $s(t)$ is the intraday periodic trend at time t of the trading day and \bar{t} is the normalized intraday time trend which is defined as the ratio of the number of seconds from opening until t to the length of the trading day in seconds. By this definition, $\bar{t} \in [0,1]$. The Appendix B of Andersen and Bollerslev (1997) presents a detailed discussion with respect to the estimation of (3.51).

In principle, the conditional expectation ψ_i and intraday periodicity component $s(t_i)$ in (3.49) and (3.50) are estimated jointly by maximum likelihood estimation. Nevertheless, when studying intraday data with ultra-high frequency (e.g., tick-by-tick transaction data), numerical difficulty rises inevitably. In order to achieve convergence for a joint estimation of ψ_i and $s(t_i)$, it is often the case that the algorithm costs a relatively long time. Due to this reason, a two-step procedure is commonly applied to simplify the estimation. That is, one first estimates the intraday periodicity component $s(t_i)$ and use the seasonal normalized \tilde{w}_i durations to estimate the ACD models. As discussed previously, the two-step estimates are consistent but not efficient. The asymptotic property of the two-step estimator can be found in Engle (2000) and Engle (2002). The derivation of the asymptotic property is based on the results of Newey and McFadden (1994) for the GMM (generalized methods of moments) estimation. For a large sample, the joint estimation and two-step estimation usually give very similar results (Engle and Russell 1998 and Bauwens and Giot 2000) as expected.

Veredas et al (2001) applied a different strategy to estimate the intraday pattern. They proposed a semiparametric estimator where the intraday periodicity is jointly estimated nonparametrically with ACD specifications for the duration dynamics. Specifically, they introduced a joint estimation of \tilde{w}_i and $s(t_i)$ in Equation (3.49) where \tilde{w}_i is modelled by the ACD specification and meanwhile $s(t_i)$ is unspecified and therefore is estimated nonparametrically.

There are also many alternative techniques used to capture the intraday pattern. Tsay (2002) applied quadratic functions with indicator variables that identify the timing of the durations.

Drost and Werker (2004) chose only one indicator variable that identifies the lunchtime for the model of Tsay (2002). The main reason is that, for their data, trading intensity are almost constant except for the durations near the lunchtime.

Dufour and Engle (2000) introduced dummy variables for the intraday pattern into their vector autoregressive (VAR) system. Interestingly, they found little evidence for the intraday pattern except that the first thirty minutes of the trading day have significantly different dynamics from the rest trading hours.

It is noticeable that, although the practice of filtering raw durations by intraday pattern is widely accepted, the effectiveness of these techniques can only be roughly examined by some stylized facts such as the U-shape trading pattern. Moreover, it also complicates the diagnostic of the ACD specifications. Concerns regarding this issue can be found in Bauwens et al. (2004) and Meitz and Terasvirta (2006). Indeed, as argued by Bollerslev (1997), given the lack of economic theory that could guide a plausible parametric form of the intraday periodicity patterns, it seems natural to estimate them by nonparametric methods.

3.2.2.2 Testing the ACD

In this section, I briefly discuss the procedures to test the ACD models. Given the vast different tests that have been developed since the introduction of ACD models by Engle and Russell (1998), an exhaustive review is beyond the scope of this Chapter. Therefore, I focus on three different categories of these tests: residual diagnostics, density forecast evaluations and Lagrange multiplier tests.

The simplest and most direct way to evaluate the goodness of fit for the ACD models is to analyze the distributional and dynamic property of the residuals. Recall the standard specification of the ACD model from (3.7) to (3.9),

$$\begin{aligned} E(w_i | \bar{w}_{i-1}, \bar{z}_{i-1}) &\equiv \psi_i, \\ \psi_i &= \psi(\bar{w}_{i-1}; \theta_\psi), \end{aligned}$$

$$w_i = \psi_i \epsilon_i.$$

Thus, the residuals are given by

$$\hat{\epsilon}_i = \frac{w_i}{\hat{\psi}_i}, i = 1, 2, 3, \dots, N, \quad (3.52)$$

where $\hat{\psi}_i$'s are the estimates for the conditional expectations ψ_i 's under the ACD specification. Suppose the specification is correct, the series $\{\hat{\epsilon}_i\}$ should be iid clearly.

Thus, one can use Ljung-Box statistics based on the centered residuals to examine whether the specification can fully capture the intertemporal dependence in the duration process.

Moreover, if the distributional assumption for ϵ_i under the ACD specification is correct, the residuals $\{\hat{\epsilon}_i\}$ should follow the distribution that corresponds to the assumption with unit mean. Thus, graphic checks such as the Quantile-Quantile plot and general goodness of fit statistics such as the EDF statistics discussed in Chapter 1 can be applied to examine the residuals.

Alternatively, one can also use moment conditions corresponding to the specified distribution to evaluate the goodness of fit. For instance, if ϵ_i is assumed to follow the standard exponential distribution, then the variance and mean of $\{\hat{\epsilon}_i\}$ should be nearly equivalent to each other. Engle and Russell (1998) proposed the statistics $\sqrt{n}(\hat{\sigma}_{\hat{\epsilon}}^2 - 1)/\sigma_{\epsilon}$, where $\hat{\sigma}_{\hat{\epsilon}}^2$ is the sample variance of the residuals and σ_{ϵ} is the standard deviation of the random variable $(\epsilon_i^2 - 1)$. For an exponentially distributed ϵ_i , $\hat{\sigma}_{\hat{\epsilon}}^2$ should be very close to 1 and the value of σ_{ϵ} is $2\sqrt{2}$. They also provided the asymptotic property of this test statistics.

Although the residuals check provides very illustrative evidence of the goodness of fit, it is also possible to examine the ACD specification by evaluating the in-sample density forecasts. Diebold et al. (1998) provided a framework of evaluating the goodness of fit based on the probability integral transform

$$q_i = \int_{-\infty}^{x_i} f_i(u) du, \quad (3.53)$$

where $f_i(x_i)$ is the sequence of one-step ahead probability density forecasts and x_i is the corresponding random process.

They showed that, if the model specification is correct, then q_i should be uniformly distributed and should be iid. Therefore, one can test the series of q_i against the uniform distribution to evaluate the in-sample density forecasts. Specifically, a goodness of fit test can be conducted by categorizing the probability integral transform q_i and calculating a chi-squared statistic that is based on the frequencies of the different categories. The chi-squared statistic can be presented as following,

$$\chi^2 = \sum_{i=1}^T \frac{(n_i - n\hat{p}_i)^2}{n\hat{p}_i}, \quad (3.54)$$

where T is the total number of categories, n_i is the number of observations in category i and \hat{p}_i is the estimated probability to observe a realization of q_i in the category i . Further details can be found in Bauwens et al. (2000) and Dufour and Engle (2000).

It is often the case that one not only is interested in testing whether the model specification is correct or not but also is interested in, when rejected, identifying the source of misspecification. In the case of testing ACD specifications, there are two main sources of possible misspecifications. The first one is the distributional assumption and the second one is the conditional expectation specification. Moreover, as discussed previously, the quasi-maximum likelihood estimation of the ACD model relies heavily on correct specification of the conditional expectation. For research that focus on the comparison of different specifications of the conditional expectation ψ_i , the validity of the conditional expectation specification might be more important than the correctness of the complete density.

A common procedure in econometric literatures for detecting the model misspecification is the Lagrange multiplier (LM) tests. A detailed discussion regarding the LM tests can be found in Engle (1984). I now present a brief discussion of the LM test for the ACD specifications. In order to conduct the LM test, one need to specify a more general model that can nest the specification of the null hypothesis. Assume that the ACD specification of the null hypothesis H_0 is a special case of the following general specification,

$$\begin{aligned} w_i &= \psi_i \epsilon_i, \\ \psi_i &= \psi_{0,i} + \theta_a' z_{ai}, \end{aligned} \quad (3.55)$$

where $\psi_{0,i}$ denotes the conditional expectation function under the null hypothesis H_0 depending on the parameter vector θ_0 , z_{ai} is the vector of missing variables and θ_a is the vector of additional parameters. The prime indicates inner product. Now, the null hypothesis H_0 that the specification is correct can be presented as $H_0: \theta_a = 0$. Consider the quasi-maximum likelihood estimation of the ACD model given by (3.24),

$$L(\theta) = \sum_{i=1}^n l_i(\theta) = - \sum_{i=1}^n \left[\frac{w_i}{\psi_i} + \ln \psi_i \right].$$

Let $\hat{\theta}_0$ be the quasi-maximum likelihood estimate under H_0 . Then, the LM test can be obtained (Engle 1984) by

$$\gamma_{LM} = l' s(\hat{\theta}_0) I(\hat{\theta}_0)^{-1} s(\hat{\theta}_0)' l \sim \chi^2(M_a), \quad (3.56)$$

where l is a vector of unit with length n , $s(\hat{\theta}_0)$ is the matrix of size $n \times (M_0 + M_a)$, and $I(\hat{\theta}_0)$ is the information matrix evaluated at $\hat{\theta}_0$. With respect to the matrix $s(\hat{\theta}_0)$, M_0 and M_a denote the number of parameters in θ_0 and θ_a respectively. Moreover, $s_{i,j}(\hat{\theta}_0) = \frac{\partial l_i(\hat{\theta}_0)}{\partial \theta_j}$ where $\theta_{j>0}$ is the j th element in $\theta = (\theta_0, \theta_a)$. In other words, the element $s_{i,j}(\hat{\theta}_0)$ in the matrix $s(\hat{\theta}_0)$ is the contribution of $\frac{\partial l_i(\hat{\theta}_0)}{\partial \theta_j}$ to the score function evaluated under H_0 .

The information matrix $I(\hat{\theta}_0)$ can be estimated according to outer product of gradients. Hautsch (2004) and Pacurar (2006) provided detailed discussion regarding the LM tests.

Along with the introduction of ACD models, Engle and Russell (1998) also proposed a method to investigate the nonlinear dependencies between residuals and the past information set. The residuals are divided into bins ranging from zero to infinity. Then, they regress the residuals on indicators showing whether the previous duration has been in one of those bins. If the residuals are iid as expected, then the regression should present no predictability and therefore the coefficients of each indicator should be zero. When rejected, the bins corresponding to indicators with significant coefficients now provide information regarding the source of misspecification.

As discussed previously, there are many literatures regarding the test of ACD models. Many of the ACD tests are developed to test specific market microstructure hypotheses. The selection of those tests depends largely on the specific question that one wants to study.

For instance, the residual diagnostics can provide valuable information for choosing the appropriate distribution assumption of ϵ_i . Meanwhile, if one wants to test the correctness of the specification of the conditional expectation ψ_i , then LM tests might be needed to give a complete picture of the goodness of fit.

3.2.2.3 ACD Papers

As discussed previously, in order to test different market microstructure hypotheses, the definition of durations is often adjusted. Thus, I categorize the empirical results of ACD literatures with respect to the definition of durations. I discuss some stylized facts that are widely identified for the durations and review the microstructure hypotheses that are often tested using ACD specifications. In general, there are two main topics that most of the researchers focus on: the model specification that can fit the dynamics of durations and the hypothesis test of different microstructure theories.

3.2.2.3.1 Trade durations

The durations can be naturally defined as the time difference between consecutive transactions as in the point process modelling of arrivals of transactions. I now call durations defined in this sense as the “trade duration”. There is, of course, a wide range of trade durations literatures. See., e.g., Engle and Russell (1998), Engle (2000), Zhang et al. (2001), Bauwens and Veredas (2004), Bauwens (2006) and Manganelli (2005).

Several stylized facts of trade durations are widely identified using different trade duration data from different markets. First, the trade durations, from different markets, are found to exhibit the phenomenon of “duration clustering” which can be viewed an analog of the “volatility clustering” for the duration series. For example, the durations have significant positive autocorrelations. In other words, long (short) durations tend to be followed by long (short) durations. Besides, the autocorrelation function decays very slowly, which in turn suggests that the persistence should be considered when modelling the trade duration dynamics. In general, the effect of the duration clustering is analyzed by the correlogram of the duration series and the LjungBox statistics. A very slowly decaying autocorrelation

functions might be indicative of the existence of long-memory characteristics of the duration process. Empirical evidence for the long memory characteristics of duration series can be found in Engle and Russell (1998), Jasiak (1998) and Bauwens et al. (2004). In fact, Jasiak (1998) developed the Fractional Integrated ACD (FIACD) model to capture this effect.

Second, most of the trade durations identify the phenomenon of “overdispersion”. That is, the standard deviation of the trade duration exceeds the mean of trade durations. The dispersion test of Engle and Russell (1998) can be applied to test this effect formally. Or simply, Dufour and Engle (2000) used a Wald test to examine the equality between the standard deviation of the trade duration and the mean of trade durations. In general, the empirical distribution of the trade durations has a hump at very short durations and long right tail (see, e.g., Bauwens et al. 2004, Giot 2001 and Engle and Russell 1998). This feature usually is interpreted as the evidence that the exponential distribution is not appropriate for the unconditional modelling of trade durations (for the standard exponential distribution, the first and second moments are equal to unit). Nevertheless, it does not necessarily mean that the conditional trade durations are not exponentially distributed.

Third, a large proportion of the trade durations sample has a value of zero or nearly indistinguishable from zero. This phenomenon suggests that there are many transactions occurring simultaneously. Moreover, the prices associated with simultaneous transactions are nearly unchanged. As discussed in Chapter 1, the raw tick-by-tick trade durations are subject to institutional features such as the data recording mechanism. Unfortunately, those are of no common knowledge.

A common approach that is used to eliminate the zeros in the trade duration series is to aggregate these simultaneous transactions. Price are averaged according to volumes. The market microstructure theory behind the procedure is split-transaction strategy of the specialists. That is, large orders are carried out by splitting into small orders. When transactions in the data do not occur simultaneously but are spaced in time with varying short intervals, the identification of split-transactions might be very difficult. Besides,

multiple transactions occurring simultaneously might be informative since it reflects a rapid pace of the market. This is line with the findings of empirical researches that the trading intensity have very strong explanatory power over the price dynamics (Ane and Geman 2000, Easley et al. 1995 and Jones et al. 1994).

Interestingly, Zhang et al. (2001) found that the exact number of the transactions occurring simultaneously has very limited explanatory power over future transactions rate, and meanwhile the existence of such multiple transactions does have significant explanatory power over future transaction dynamics. They further incorporated this effect into their TACD model by introducing a lagged indicator representing the existence of simultaneous transactions into the specification of conditional expectation of trade durations. Another explanation for zero trade durations is provided by Veredes et al. (2001). They argued that simultaneous transactions are possibly caused by the efforts of traders to trade at the round prices. Specifically, traders tend to post limit orders to be executed at the round prices. Their theory is based on the fact that the prices of simultaneous transactions are clustered around round prices.

With respect to the modelling of trade durations, the most frequently used model is the standard ACD and the log-ACD with low orders of the lagged term such as the ACD (1,1) or log-ACD (1,1) (Engle and Russell 1998, Engle 2000, Zhang et al. 2000 and Fernandes and Gramming 2006). Interestingly, most of the authors found that the residuals still exhibit significant autocorrelations. This phenomenon is viewed as a signal of the nonlinear dependence in the duration series because the standard ACD models are linear. Bauwens et al. (2004) provided an exhaustive review and reported that almost all existing ACD models (standard ACD, log-ACD, TACD, EACD among the others with various distributional assumptions) cannot provide a satisfied modelling of the conditional expectation of durations. Besides, Dufour and Engle (2000), Engle and Russell (1998) noticed that the ACD specification with the best in-sample test results generally does not provide the best out-sample forecasts. The persistence of the ACD specification are found to be very high (Engle and Russell 1998, Engle 2000 and Dufour and Engle 2000). For example, for a standard ACD model the sum of the two parameters α and β is very close

to one.

Regarding testing market microstructure hypothesis using trade durations, Engle (2000) used UHF-GARCH model to investigate the relationship between volatilities and trade durations. He found a significant negative correlation between conditional expectation of trade durations and volatility, which is in line with the theory of Easley and O'Hara (1992) where "no trade means no news". Besides, they found that, when adding the lagged bid-ask spread as an explanatory variable into the conditional volatility specification, the coefficient of the spread is positive. In other words, a high spread is sign of rising volatility. Similar results can also be found in Feng et al. (2004), where they also found a negative correlation between the trade durations and realized volatility.

In contrast, Gramming and Wellner (2002) studied the interdependence between trading intensity and volatility. They found that the trading intensity are negatively correlated with lagged volatility, which is consistent with the Admati and Pfleiderer (1988). According to their results, the trade durations are positively correlated with lagged volatility. Engle and Russell (2005) applied an ACM-ACD model on the tick-by-tick data of Airgas on NYSE and found that the long durations are associated with declining prices. This finding is in line with the theory of Diamond and Verrechia (1987), where "no news is the bad news".

There are several remarks that be summarized by the above discussion with respect to the trade duration modelling.

First, the empirical results suggest that, while there are a wide range of ACD specifications, their performance in the modelling of conditional expectation of trade durations are far away from satisfactory. The uncertainty of misspecification further complicates this issue. It naturally rises the question to select the most appropriate distributional assumptions.

Second, most of the researches using data from NYSE, usually blue chips with very high liquid. The transaction record is subject to institutional features that are hard to analyze and eliminate.

Finally, while in general the empirical results regarding the interdependence of volatility and durations support the theory of Easley and O'Hara (1992), partially contradictory empirical evidence can be found in a way similar to the theoretical microstructural models.

3.2.2.3.2 Price durations

As discussed in Section 3.2.1.3.2 with respect to the relationship between duration and volatility, sometimes selecting procedures are applied to the raw tick-by-tick duration series to test some market microstructure hypotheses.

In general, a large proportion of the tick-by-tick trade durations are associated with unchanged prices. Thus, sometimes special attentions are paid to transactions that are associated with certain change in the prices. That is, given the raw tick-by-tick transaction data, one keeps the transactions with certain price movements, the time difference between those transactions are defined correspondingly as the durations.

Specifically, consider certain amount of price movement, C . Suppose the sequence of tick-by-tick transaction arrival time is $\{t_0, t_1, \dots, t_n\}$ and the associated price series is $\{p_{t_0}, p_{t_1}, \dots, p_{t_n}\}$. Let $\tau_0 = t_0$, the series $\{\tau_0, \tau_1, \dots, \tau_m\}$ are defined inductively by $\tau_i = t_s$ where $s = \min\{j \mid |p_{t_j} - p_{\tau_i}| \geq C, t_j > \tau_i, j = 1, 2, 3, \dots, n\}$ for $i > 0$. For instance, since $\tau_0 = t_0$, τ_1 would be the first t_j such that $|p_{t_j} - p_{t_0}| \geq C$ and τ_2 would be the first t_j such that $|p_{t_j} - p_{\tau_1}| \geq C, t_j > \tau_1$. I will call the duration series in which each duration is associated with certain price changes as the “price durations”. One can immediately infer from the definition that the price duration is closely related with the volatility since each duration is associated with a price movement that is equal to or larger than C . The conditional expectation ψ_i of price durations then can be roughly understood as an indicator of the time needed to expect a price movement larger than C .

Engle and Russell (1998) presented a clear parametric relation between the price duration and the condition instantaneous volatility. Specifically, let the information set until i th transaction be F_i . Then the conditional instantaneous intraday volatility is related with the

conditional hazard of the price durations by

$$\sigma^2(t|F_{i-1}) = \left(\frac{c}{p(t)}\right)^2 h(w_i|F_{i-1}), \quad (3.57)$$

where $\sigma(t|F_{i-1})$ is the conditional instantaneous volatility, $p(t)$ is the price and $h(w_i|F_{i-1})$ is the conditional hazard function.

In general, empirical research finds that the price durations share many similar characteristics with the trade durations such as “duration clustering”, overdispersion, right-skewed empirical distribution (Engle and Russell 1998, Bauwens and Giot 2000, Bauwens et al. 2004 and Fernandes and Gramming 2006). Interestingly, under the ACD specifications, compared with trade durations, the residuals of price durations are much more regular with respect to the serial dependence. Empirical evidence regarding this phenomenon can be found in Bauwens et al. (2004) and Fernandes and Gramming (2006).

With respect to the model specification, Bauwens et al. (2004) argued that the simple ACD model such as the standard ACD and log-ACD outperforms the complicated models such as the TACD and SVD. They also suggested that, compared with the model specification, the distributional assumption seems to have stronger influence on the forecast performance of the ACD models.

As discussed previously, the definition of price durations itself can be viewed as a measurement of the volatility. Therefore, the ACD modelling for price durations can be viewed as a volatility modelling method.

Papers with respect to this consideration can be found in Giot (2000), Gerhard and Hautsch (2002) and Giot (2002). For instance, Giot (2000) applied an ACD specification to the price durations of IBM and directly used the estimated ACD conditional expectation function to compute the intraday volatility by (3.57). Giot (2002) further built an Intraday value at risk model where the volatilities are computed according to the price durations modelled by log-ACD specifications. However, empirical results suggest that the price duration based model fails most of the time for all stocks under consideration. Giot (2002) argued that the

unsatisfactory results might be due to the normal assumption imposed on intraday returns.

With respect to the market microstructural hypothesis tests, the empirical results from price durations are in general very similar to the results from trade durations. For instance, Engle and Russell (1998) and Bauwens and Giot (2000) found out that the price durations are negatively correlated with volatility, trading intensity and volumes.

3.2.2.3.3 Volume durations and others

It is clear that the tick-by-tick trade durations can be thinning by volumes in the same way as by prices. It was first introduced by Jasiak et al. (1999) as a measure of liquidity since it measures the time needed for given amount of volume to be traded. In this sense, it could be understood as the time cost of liquidity.

Compared with the trade durations and price durations, there are much fewer research applying ACD models to volume durations (Bauwens et al. 2004, Bauwens and Veredas 2004 and Fernandes and Gramming 2006). The empirical results suggest that the volume durations have statistical characteristics different from the price durations and trade durations. The standard deviation of volume durations tends to be smaller than the mean, suggesting the “under-dispersion”. Meanwhile, as discussed previously, the trade durations and price durations exhibit overdispersion.

Regarding the ACD specification, Bauwens et al. (2004) suggested the use of standard ACD and log-ACD with Burr innovations, which means that the ϵ_i in (3.9) is assumed to have a Burr density. He also argued that, as in the case of price duration and trade duration, sophisticated ACD model cannot outperform the simple ACD models with respect to the predictability.

The ACD modelling of irregularly spaced financial data is not limited to the intraday trading process. Interesting applications of ACD models to the risk management can be found in Christoffersen and Pelletier (2004) and Focardi and Fabozzi (2005). For instance, Christoffersen and Pelletier (2004) used it to back test the Value at Risk models.

Specifically, they applied ACD specifications to model the time difference between violations of some VaR models. If the VaR is correctly specified, then the conditional expectation of durations for the violations should be a constant that is decided by the significance level of the VaR model.

3.3 Data and some empirical results

In this section I present the statistical description of the data base. In Section 3.3.1 I present a discussion of the market microstructural issue of “split-transactions” associated with the tick-by-tick trade durations of SPY on NYSE. I further introduce the durations associated with the transaction-aggregated returns that is introduced in Chapter 1 and Chapter 2. In Section 3.3.2 I present the statistical descriptions of the transaction-aggregated durations. Section 3.3.3 includes the analysis of the intraday periodicity of durations. Section 3.3.4 presents the unconditional distributional properties of the durations.

3.3.1 Market microstructural issues and transaction-aggregated durations

As introduced in Chapter 1 and Chapter 2, my data is the tick-by-tick transaction records of the symbol SPY on NYSE. The data ranges from 01/02/2014 to 12/31/2014. There are 252 trading days during this period. The out-of-hour transactions whose occurrence lie outside the time period between 9:30 am and 4:00 pm on each day are removed. There are 79156264 tick-by-tick observations in the sample. Three random variables are associated with each observation in the tick-by-tick transaction records: timing of transaction, execution price of transaction and trading volumes.

Before my presentation of the tick-by-tick durations, there are several market microstructural issues that I will discuss. These issues might be considered as irrelevant for studies using sample from fixed time interval. However, for a tick-by-tick transaction data of a symbol with ultra-high liquidity like SPY, the interpretation of the data with respect to transaction arrivals depends largely on the understanding of those issues.

The first issue that may arise when using tick-by-tick transaction data is the difficulty of

matching trades and quotes. This problem occurs if the trades and quotes are separately recorded and stored like the Trade and Quote (TAQ) database released by NYSE. Fortunately, my data is not directly quoted from the TAQ. The tick-by-tick transaction record is obtained from the commercially available data company of Tick Data and the trades and quotes are matched automatically. Nevertheless, this convenience comes at the cost of the information of bid and ask prices.

The second issue is the so called “split-transactions” effect. This phenomenon has been mentioned in the previous discussion regarding price durations. I now elaborate it and explain my methodology of dealing with this issue. Split-transaction occurs when an order (usually of large size) on one side is matched with several smaller orders on the opposite side. An immediate result is that the time difference between these transactions is extremely small. Correspondingly, the execution prices of these transactions are equal or are monotonic (increasing or decreasing, depending on how the large order is matched). For an electronic trading system, the time difference between these transactions is usually determined by the measurement accuracy. Such recording mechanism also decides whether those transactions will be considered as simultaneously occurring. In most of the research, this issue is solved by aggregating transactions occurring at the same time. However, if these transactions are carried out with extremely small time differences as in my data, the identification of split-transactions will be very difficult. Moreover, since the bid-ask quotes are not available for our data, I cannot identify the possible split-transactions according to bid-ask information.

In Chapter 1 and Chapter 2, I introduced the concept of transaction-aggregated returns. Compared with time-aggregated returns, the transaction-aggregated returns exhibit stronger aggregational normality. Detailed discussion is presented in section 1.3. More importantly, it preserves the information contained in the time of transactions. Moreover, the durations between each of observation in the transaction-aggregated returns contain equal number of transactions. Thus, those durations can be considered as a direct measure of trading intensity. Given the widely recognized relation between trading intensity and volatility dynamics, it seems valuable to investigate the durations defined in this sense. I

will use the terminology “*transaction-aggregated durations*” to refer to the durations that correspond to the transaction-aggregated returns defined in Chapter 1.

Specifically, let $p_{d,n}$ be the execution price of the n th transaction at day d and let $t_{d,n}$ be the associated time of transaction. Therefore, the tick-by-tick trade duration is $w_{d,n}$, where $w_{d,n} = t_{d,n} - t_{d,n-1}$ (denote the time of the first transaction at day d as $t_{d,0}$).

The n th tick-by-tick return at day d is, $r_{d,n} = \ln \left(\frac{p_{d,n}}{p_{d,n-1}} \right)$. Suppose we use every consecutive T tick transactions to calculate our non-overlap transaction-aggregated return series. Then $(T)r_{d,i}$, the i th transaction-aggregated return on day d , is defined as

$$(T)r_{d,i} = \ln \left(\frac{p_{d,(T-1)i}}{p_{d,(T-1)(i-1)}} \right).$$

Correspondingly,

$$(T)w_{d,i} = t_{d,(T-1)i} - t_{d,(T-1)(i-1)}.$$

Thus, the tick-by-tick trade duration can be viewed as the special case where $T=2$. Moreover, although the duration between consecutive transactions has been widely considered as a random variable, its distributional property is insufficiently explored. This might be caused by the difficulty arising with the previously discussed market microstructural issues. The definition of transaction-aggregated duration is based on the sum of tick-by-tick durations. Therefore, study with respect to the aggregational property of tick-by-tick durations can provide valuable knowledge regarding the nature of the tick-by-tick durations. This is very similar to the case of intraday returns where the aggregational normality of intraday returns in turn provides support for the stable assumption on returns.

In fact, given a tick-by-tick transaction records, the sequence of tick-by-tick transaction time naturally forms a partition for the time interval between the first and last tick-transaction. Once a definition of duration is chosen, it actually just retains some of the tick-by-tick transactions and therefore a new partition of the time interval between the first and last transaction is formed. The tick-by-tick trade durations is a refinement of the new duration sequences. In this sense, the traditional use of fixed time interval in the analysis of a financial time series is equivalent to imposing a regular partition of the time.

3.3.2 The descriptive statistics of transaction-aggregated durations

I first present the descriptive statistics of the transaction-aggregated durations. As in Chapter 1 and Chapter 2, let T represent the number of tick-by-tick transactions that are used to calculate the transaction-aggregated durations. For instance, $T=400$ indicates that the durations are now the time difference between every four hundred tick-by-tick transactions. The values of T are determined by the sample size of time-aggregated transaction series based on some fixed time interval. Specifically, suppose I measure the return per one second from the tick-by-tick transaction data. There will be $390 \times 60 \times 252 = 5896800$ observations in the 1-second intraday return series. Correspondingly, it means in average each 1-second return is generated by $\frac{79156264}{5896800} \approx 13.42$ tick-by-tick transactions. In order to explore the transaction-aggregated returns at high frequency level, T is chosen to be 13 and 67 to correspond to 1s and 5s time intervals.

	Sample size	Mean	Standard Deviation	q (0.05)	q (0.25)	q (0.5)	q (0.75)	q (0.95)	Ljung-Box statistics
Tick durations	79156012	0.074495162	0.459798842	0	0	0	0.005	0.423	1974424
T-13 durations	6596215	0.893933316	2.601303677	0.001	0.006	0.196	1.014	4.035	342587
T-67 durations	1199206	4.914029642	9.170544074	0.047	1.17	3.018	6.432	16.187	85805
T-134 durations	595032	9.900267661	14.67155606	0.693	3.252	6.757	12.985	29.676	73310
T-400 durations	198261	29.7117307	34.47228761	4.579	12.164	21.995	39.197	80.22945	42351
T-800 durations	98942	59.33634262	47.88790054	11.046	26.195	45.426	78.793	153.4252	71763
T-1200 durations	65893	89.08588393	69.41574311	17.73315	40.68475	69.259	118.2125	225.13925	48350
T-2400 durations	32876	178.4973975	131.302421	39.4892	85.379	141.88	237.111	435.7287	20892

Table 3.1 Descriptive statistics of transaction-aggregated durations

The $q(p)$ in Table 3.1 represents the quantile of the sample corresponding to the probability p , Ljung-Box Q statistics is calculated for the first 20 lags of the sample autocorrelation function. I keep durations with length of zero to give a complete picture of the raw tick-by-tick data.

There are several distinctive characteristics of the tick-by-tick durations of SPY that can be concluded from Table 3.1. First, the SPY has high liquidity. The mean of the tick-by-tick durations is around 0.0745s. Further evidence can be found in the quantiles. The 0.5-quantile of the tick-by-tick durations is zero and the 0.95-quantile is only 0.423s. The quantiles are directly calculated from the empirical distribution. For instance, the 0.95-quantile is solved from the equation $F(x) = \Pr(X \leq x) = 0.95$, where $F(x)$ is the empirical distribution function. In this case, the microstructure issues discussed previously have very strong impact on the pattern of tick-by-tick durations since the tick-by-tick durations have very small average that is very close to zero. The split-transaction effect is strong.

Second, the tick-by-tick durations exhibit strong overdispersion since the standard deviation exceeds the mean significantly. Given a mean of 0.0745 the standard deviation of the tick-by-tick standard deviation is 0.4598. This can be viewed as an evidence against the exponential distribution since the exponential distribution has a mean equal to its standard deviation.

Finally, as suggested by the Ljung-Box statistics, the null hypothesis that the first 20 lags in the autocorrelation functions have coefficients zero is clearly rejected. In contrast, the tick-by-tick durations have very strong autocorrelation, suggesting “duration clustering” effect.

In general, the transaction-aggregated durations exhibit characteristics similar to the tick-by-tick durations. For instance, the Ljung-Box statistics are of large values and the standard deviations exceed the mean in general. Nevertheless, it is noticeable that, as the number used to aggregate the tick-by-tick duration increases, the phenomenon of overdispersion gradually dies out. Specifically, the standard deviation of the $T=800$ durations is 47.89, and the mean is 59.34.

Figure 3.1 presents the autocorrelation function of the tick-by-tick durations. As expected, I observe highly significant autocorrelations. The first lag autocorrelation is around 0.2 and

the autocorrelations decay with the lags slowly for the first 50 lags. It then almost remains stable for the rest of the lags. It suggests the tick-by-tick duration process is very persistent. In order to explicitly present the decaying autocorrelation structure of the transaction-aggregated durations, I exhibit the correlogram of $T-13$ and $T-67$ durations for up to 500 lags. Note that since these two duration series are aggregated from the tick-by-tick duration series, now each lag represents longer time interval.

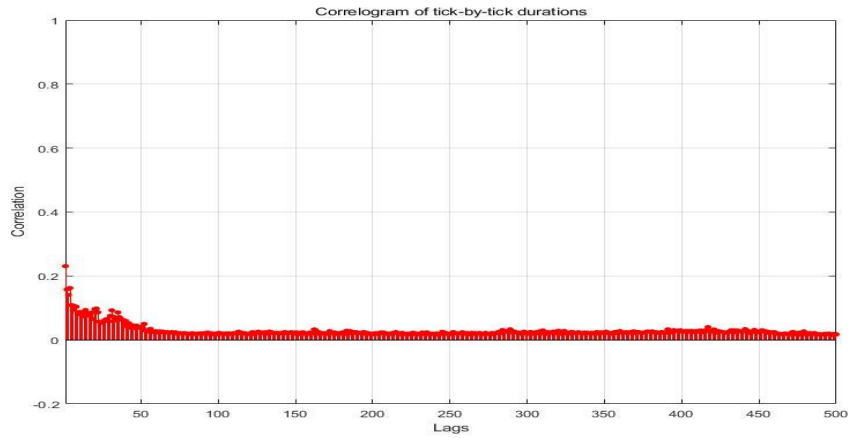


Figure 3.1 Correlogram of tick-by-tick durations




Figure 3.2 Correlogram of T-13 durations




Figure 3.3 Correlogram of T-67 durations

Figures 3.2 and 3.3 present the correlogram of the $T-13$ and $T-67$ durations. The findings are in line with Figure 3.1. Observe the autocorrelations now decay with the lags at a slow, almost hyperbolic rate that is typical for long memory process. Besides, the $T-13$ durations have a very strong first lag autocorrelation around 0.39.




Figure 3.4 The tick-by-tick duration of SPY

Figure 3.4 presents the tick-by-tick duration of SPY for the whole sample. The y-axis is limited to 50s. Although the tick-by-tick durations range from zero to almost 1800 (in seconds), most of the tick durations are shorter than 1s. Besides, as suggested by figure 3.4, the tick durations exhibit high variability.

3.3.3 Intraday seasonality of durations

It is widely assumed that intraday transactions present very strong periodic pattern over the

trading day. The discussion with respect to the intraday volatility periodicity of SPY can be found in Chapter 2. Consequently, before applying the ACD model to the durations of SPY, it is necessary to investigate the intraday seasonality of durations first.

In order to illustrate the intraday pattern of tick-by-tick durations, I first present the plot of average durations for 1-minute intervals over the trading day. Specifically, given a specific trading day, I count the number of tick-by-tick transactions falling in each of the 1-minute intervals and use the averages of tick durations in the 1-minute intervals to represent the tick duration dynamics over the trading day. Based on the empirical literatures and results from Chapter 2, I should expect an inverted U-shape pattern of the tick durations over the trading day since the opening and closing trading hours are associated with high trading intensity. Moreover, as suggesting by Table 3.1 and Figure 3.4, the durations should be small (very close to zero) except for the “lunch time”.




Figure 3.5 The average tick-by-tick duration of SPY over the trading day

Figure 3.5 presents an inverted U-shape of the tick-by-tick durations over the trading day. The average tick duration ranges from a low of 0.025s at the opening and closing of market to a high of around 0.35s at the middle day. Specifically, the tick-by-tick duration starts around 0.025s and gradually increases with the 1-minute lags. It reaches the peak around the 250th 1-minute lag which corresponds to the calendar time of 13:40. Then it decays slowly to 0.025 at the closing of the market. An interesting finding is that, compared with the tick durations in the morning, the tick durations in the afternoon exhibit much higher

variability (observe the dramatic fluctuation of tick durations after the 200th 1-minute lag). As expected, the tick durations are of very small values even in the middle day.




Figure 3.6 The correlogram of SPY average tick-by-tick durations

In order to investigate the dynamic feature of the average tick-by-tick durations, I further plot the correlogram of the average tick-by-tick durations for up to five days. The plot is exhibited in Figure 3.6. The question that I want to address is, whether the pattern depicted in Figure 3.5 cycles at daily basis.

Figure 3.6 gives a very surprising answer with respect to the intraday periodicity of durations. If the pattern depicted in Figure 3.5 is periodic at daily frequency, we should observe a corresponding U-shape in the correlogram as in the case of intraday periodicity of volatility (Figures 2.5 and 2.6 of Chapter 2). In contrast, Figure 3.6 suggests the autocorrelations of average tick-by-tick durations per minute decays monotonically with the 1-minute lags. The autocorrelations become insignificant for lags of order higher than 180. In other words, the average tick duration in the current minute is uncorrelated with the average tick duration three hours later.

This result is not necessarily contradictory to Figures 3.1, 3.2 and 3.3. Indeed, as suggested by those three figures, the autocorrelations of the transaction-aggregated durations (including the tick-by-tick durations where $T=2$) decay very slowly with lags. However, their autocorrelations do not necessarily decay slowly with time. Due to the different time interval that each duration represents, I cannot directly identify how the autocorrelation of the transaction-aggregated durations decay with time. But we can gauge this dependence

by the mean of the transaction-aggregated durations. For instance, the mean of the tick-by-tick duration series is around 0.0745s. Then the time interval between the 1st and 500th lag in Figure 3.1 can be very roughly interpreted as $0.0745\text{s} \times 500 = 37.245\text{s}$. Meanwhile, the average tick-by-tick durations in Figure 3.6 represents the average value of tick-by-tick durations in each 1-minute interval. Thus, the time difference between two consecutive lags is one minute.

The characteristics of the average tick-by-tick duration per minute can be very different from the tick-by-tick durations. In order to explicitly explore the dynamic feature of the tick-by-tick durations, one might be interested in the correlogram of the raw tick-by-tick durations for lags of higher orders. However, given the extremely small mean of tick durations (0.0745s), if I want to explore the autocorrelation structure of tick-by-tick durations for detecting some daily effect, e.g., whether the intraday pattern of Figure 3.5 is periodic on daily basis, an extremely large number of lags will be required. Given the tick-by-tick sample size of 79156012, the computational difficulty arises due to the high requirement of random-access memory.

Moreover, the pattern of tick-by-tick durations of SPY depends largely on the recording mechanism as discussed previously. Unfortunately, those are not common knowledge and therefore I cannot exclude their influence using my data. Thus, I exploit the autocorrelations of the transaction-aggregated durations to circumvent those difficulties.




Figure 3.7 The correlogram of T-800 transaction-aggregated durations

Figure 3.7 presents the correlogram of the *T*-800 durations for up to 1500 lags. Given the

mean of 59.34s for the $T-800$ transaction-aggregated durations, Figure 3.7 can be roughly interpreted as the correlogram of $T-800$ durations for up to four trading days. Figure 3.7 suggests that there is a U-shape pattern in the autocorrelations of the $T-800$ durations. Although the U-shape pattern gradually dies out with the lags, it is still strikingly regular for the first 300 lags. Specifically, the distance between the first peak and the second peak of the autocorrelations is around 300 lags. Since the mean of the $T-800$ transaction-aggregated durations is around one minute, 300 lags represent roughly one trading day. Figure 3.7 implies that, although the intraday seasonality of SPY durations is not as regular as the intraday seasonality of SPY volatilities, it still has unneglectable influence on the autocorrelation structure of the durations.




Figure 3.8 The correlogram of T-400 transaction-aggregated durations

Figure 3.8 gives the correlogram of the $T-400$ transaction-aggregated durations for up to 3000 lags. The pattern of autocorrelations in Figure 3.8 is very similar to Figure 3.7.

Figures 3.7 and 3.8 suggest that a direct ACD modelling of the transaction-aggregated durations might be inappropriate since the ACD models impose a geometric decay on the duration autocorrelations. It cannot capture the autocorrelation pattern in Figures 3.7 and 3.8. Thus, it is necessary that I exclude the intraday seasonality of durations before applying ACD specification to model the transaction-aggregated durations.

As discussed in the Section 3.2.2.1, a widely accepted method that is used to exclude the intraday seasonality from the raw durations is to consider following decomposition of durations w_i , $w_i = \tilde{w}_i s(t_i)$. $s(t_i)$ is considered as the deterministic “time-of-day”

component (Engle and Russell 1998). With respect to the estimation of $s(t_i)$, the most used procedures are as follows. First, one takes averages of durations over fixed time intervals of the trading day. Second, cubic spline functions are applied to smooth the averages.

The following Figure 3.9 exhibits the cubic spline function with 15-minute nodes of SPY tick-by-tick durations.




Figure 3.9 Cubic spline function with 15-minute nodes of SPY tick-by-tick duration

Once the tick-by-tick durations are adjusted according to the cubic spline function in Figure 3.9. I can easily calculate the seasonal adjusted transaction-aggregated durations by the sums of corresponding seasonal adjusted tick-by-tick durations.

Given the complicated intraday seasonality of tick-by-tick durations, this de-seasonal procedure might be too simple. I choose this commonly accepted procedure for the following reasons. First, as suggested by Figures 3.7 and 3.8, the autocorrelations of transaction-aggregated durations indeed present some periodic feature at daily frequencies. Second, due to the lack of economic theory that could facilitate a parametric specification of the modelling of intraday seasonality of durations, it seems plausible to choose a relatively simple nonparametric method. Third, the intraday dynamics of tick-by-tick durations are subject to strong influences of the market microstructural issues such as the “split-transactions” effect that we unfortunately are unable to identify. Therefore, it is difficult to provide an accurate modelling of the intraday seasonality.

3.3.4 Unconditional distributional properties of the durations

In this section, I provide an investigation with respect to the unconditional distributional properties of the transaction-aggregated durations. I use the raw durations instead of the de-seasonalized durations in this section since the analysis mainly serves as a guidance in the selection of distribution assumptions for ACD models. Besides, the raw durations give a complete description of the sample that I am interested in.




Figure 3.10 The kernel density function of tick-by-tick durations

Figure 3.10 exhibits the kernel density function of the tick-by-tick durations of SPY. Let $\{x_0, x_1, \dots, x_n\}$ be an iid sample following the unknown density f , the kernel density estimator of f at x is $\frac{1}{nh} \sum_{i=1}^n K(\frac{x-x_i}{h})$, where K is the non-negative kernel function and $h>0$ is the smoothing parameter. I use the normal kernel $K(x) = \phi(x)$ where $\phi(x)$ is the standard normal density function. Due to the large quantity of computation required to fully exhibit the kernel density function (the maximum of tick-by-tick durations is around 1800s), I only plot the function for tick-by-tick durations less than 10s (the 0.95 quantile of the tick-by-tick sample is around 0.42s).

The kernel density function for T -134 durations, T -400 durations and T -800 durations are presented in the following Figures 3.11, 3.12 and 3.13 respectively. The range of x-axis is from 0 to 360. The kernel density functions of the transaction-aggregated durations have humps on the short durations and have relatively long right tails. Besides, the aggregational

characteristics is obvious. The kernel density function of tick-by-tick durations in Figure 3.10 has a roughly inverted S-shape and meanwhile the kernel density functions in Figures 3.11, 3.12 and 3.13 are much regular and are similar to Weibull or Gamma densities. Figures 3.11, 3.12 and 3.13 intuitively suggest that an exponential assumption on the standardized duration ϵ_i might be wrong. Weibull distribution and generalized gamma distribution seems to be more appropriate for describing transaction-aggregated durations.




Figure 3.11 The kernel density function of T-134 durations




Figure 3.12 The kernel density function of T-400 durations




Figure 3.13 The kernel density function of T-800 durations

Given the characteristics of the tick-by-tick trade duration presented in Table 3.1 and Figure 3.10, continuous distributions that are widely used in ACD duration modelling such as Weibull distributions and exponential distributions cannot provide satisfactory description of the unconditional tick-by-tick durations.

Based on the empirical evidences from Table 3.1 and Figure 3.10 - 3.13, I focus on the following two issues. First, I want to analyze the effect of the market second-by-second operational details on the tick-by-tick durations. Although I cannot directly identify the effect by our data, we can gauge it by the following method. I divide the tick-by-tick durations into different categories according to their length. Then, I study the distributional properties of durations in different categories. According to market microstructure literatures, short durations are often the results of liquidity trading and meanwhile long durations are more likely to be related with informed trading (Admati and Pfleiderer 1988). Consequently, tick-by-tick durations in different categories should have very different distributional properties.

Second, I want to explore the aggregational characteristics of the tick-by-tick durations.

With respect to the first issue, I categorize the tick-by-tick durations into the followings three categories. The first category contains tick-by-tick durations with length less than 1s. Since the tick-by-tick durations have a 0.95 quantile of 0.43s, it contains most of the tick-by-tick duration data. According to Admati and Pfleiderer (1998), tick-by-tick duration in this category should be the results of liquidity transaction and therefore the market operational details play a decisive role in determining their dynamics.

It is interesting to explore the unconditional distribution of the tick-by-tick duration under such circumstances. I remove all the zero observations in this category. The second category contains tick-by-tick durations with length between 1s and 60s and the third category contains tick-by-tick durations with length larger than 60s. The following Table 3.2 summaries the descriptive statistics of the three subsamples.

	Sample size	Mean	Standard Deviation	q (0.05)	q (0.25)	q (0.5)	q (0.75)	q (0.95)
$w \in (0, 1]$	31929162	0.099620631	0.191059825	0.001	0.001	0.009	0.095	0.559
$w \in (1, 60]$	1418818	1.899863586	1.164725964	1.02	1.197	1.535	2.163	4.008
$w \in (60, \infty]$	110	185.2366909	211.4020134	63.521	85.811	117.219	194.309	567.133

Table 3.2 Descriptive statistics of tick-by-tick durations with different length

Table 3.2 indicates some characteristics of the SPY tick-by-tick durations that are certainly determined by market operational details. First, the smallest unit of time that can be recorded by the market is 0.001s. This can be easily verified by the quantiles of durations less than 1s. As a result, for durations with extremely small values, the empirical distribution presents strong discreteness. For instance, the 0.5 quantile of the first subsample is 0.09 but the smallest difference between two observations is 0.001.

Second, a large proportion of the SPY tick-by-tick durations are less than 0.1s. The sample sizes of the durations in the three subsamples are 31929162, 1418818 and 110 respectively. Moreover, the 0.75 quantile of the durations that are less than 1s is 0.095.

Third, although the durations less than 1s and durations longer than 60s both exhibit “overdispersion” (sample mean less than sample standard deviation), the durations with length between 1s and 60s do not have this characteristic. Specifically, the durations in the second subsample have a sample mean of 1.899 and a standard deviation of 1.165.

Based on these empirical findings, I can conclude that the tick-by-tick duration of SPY is heavily influenced by the market microstructural effects such as the accuracy of recording systems and the “split-transactions”. I now further investigate the distributional properties of the tick-by-tick trade durations. I fit the exponential distribution, Weibull distribution, the gamma distribution and the generalized pareto distribution to the first two subsamples of the tick-by-tick durations. Specifically, the four continuous distributions are estimated by MLE using the tick-by-tick durations in each subsample normalized by corresponding sample mean.




Figure 3.14 The ECDF of tick-by-tick durations less than 1s

Figure 3.14 presents the plot of the empirical cumulative distribution function of durations in the first subsample and the cumulative distribution functions of the four continuous distributions estimated from the subsample. The results regarding the estimation are provided in Table 3.3. The details regarding the maximum likelihood estimation are presented in Appendix D.

Exponential Distribution		
Parameter	Standard Error	Log likelihood
mu=1	0.000176973	-3.19E+07
Weibull Distribution		
Parameter	Standard Error	Log likelihood
(scale) a=0.4069	0.000165374	-8.18E+06
(shape) b=0.4623	6.20E-05	
Gamma Distribution		
Parameter	Standard Error	Log likelihood
(scale) a=0.3316	6.61E-05	-1.03E+07
(shape) b=3.0153	1.10E-03	
Generalized Pareto Distribution		
Parameter	Standard Error	Log likelihood
(scale) sigma=0.0571	2.51E-05	-6.50E+06
(shape) k=2.0664	5.44E-04	
(location) theta=0	0	

Table 3.3 The fitting of Exponential, Weibull, Gamma and Generalized Pareto distribution using durations less than 1s

None of the four distributions can provide a satisfactory modelling of the unconditional distribution of tick-by-tick durations in the first subsample. The discreteness reported previously cannot be easily captured by the four continuous distributions that are widely used in duration modellings. Thus, ignoring such effects might be wrong for ACD duration dynamics modellings.




Figure 3.15 The ECDF of tick-by-tick durations with length between 1s and 60s

Exponential Distribution		
Parameter	Standard Error	Log likelihood
Mu=1.89986	0.00159499	-2.32 E+06
Weibull Distribution		
Parameter	Standard Error	Log likelihood
(scale) a=2.14715	0.00106326	-1.94E+06
(shape) b= 1.80038	9.36E-04	
Gamma Distribution		
Parameter	Standard Error	Log likelihood
(scale) a= 0.41307	0.00527462	-1.73274E+06
(shape) b=4.59937	0.000500552	
Generalized Pareto Distribution		
Parameter	Standard Error	Log likelihood
(scale) sigma=0.709601	0.00096356	-1.23614E+06
(shape) k=0.214296	0.00108232	
(location) theta=1	0	

Table 3.4 The fitting of Exponential, Weibull, Gamma and Generalized Pareto distribution using durations with length between 1s and 60s

Figure 3.15 and Table 3.4 summarize the fitting regarding durations in the second subsample where observations vary from 1s to 60s. For clearer presentation, the x-axis is limited to 10. It can be concluded from Figure 3.15 and Table 3.4 that the four continuous distributions provide much closer approximations for the durations in the second subsample than for the durations in the first subsample. Specifically, the fit of Gamma distribution and generalized Pareto distribution seem to be promising. Although the fitting cannot be considered as statistically successful, it still implies that the characteristics of durations depends largely on the length of durations. Consequently, it intuitively suggests the use of a regime switching ACD specification for modelling tick-by-tick duration dynamics.

For the durations longer than 60s, I present the histogram to describe its distributional characteristics since there are only 110 observations. The histogram is presented in Figure 3.16. These durations are highly likely to be associated with some market events.

Based on the empirical results of the three subsamples, I can conclude several important remarks. First, for a symbol with high liquidity like SPY, the tick-by-tick durations are often with very small values (the limit of recording accuracy is often reached) that are results of the market microstructural effects. Second, durations with different length have very different characteristics. As suggested by microstructure literatures, the long durations might convey information very different from the information carried by short durations. Consequently, when modelling tick-by-tick durations in the ACD framework, a regime-switching specification such as the model of Hujer and Vuleti (2005) might be more realistic. Finally, given the dominating rules of market operational details in determining the pattern of tick-by-tick durations, ACD models might not be very suitable for describing the arrivals of transactions if one cannot explicitly exclude the effects of these details. For instance, the raw tick-by-tick duration sample of SPY contains 79156012 observations. Meanwhile, the number of observations equal to 0, 0.001 and 0.002 are 45807922, 8597959 and 2302916 respectively. It is highly unlikely that such discreteness can be captured by ACD specifications based on continuous distributional assumptions since there will be

many jumps in the empirical distribution of durations.




Figure 3.16 The histogram of tick-by-tick durations longer than 60s

I now use Kolmogorov statistics D , the Cramér–von Mises statistics W^2 and the Andersen–Darling statistics A^2 , to formally test the unconditional distribution of transaction-aggregated durations. The procedures of applying EDF statistics can be found in the Appendix A of Chapter 1. In order to investigate the aggregational distributional property of tick-by-tick durations, I fit the exponential distribution, Weibull distribution, the Gamma distribution and the generalized Pareto distribution to transaction-aggregated durations. The first three distributions are nested in the generalized Gamma distribution and meanwhile the Weibull distribution is the limiting distribution of the Burr distribution. All zeros are removed from the sample.

Table 3.5 reports the EDF statistics for the goodness of fit for the transaction-aggregated durations whose descriptive statistics are given in Table 3.1. The distributions are estimated by MLE and Monte-Carlo simulations are applied to give critical values. Specifically, parameters of each distribution are estimated by corresponding transaction-aggregated durations and then random numbers are generated by the estimated distribution for the calculation of critical values. Details are presented in Appendix D. I do not include the critical values in Table 3.5 since the EDF statistics exceed the critical values significantly.

As suggested by the extremely high values of the three EDF statistics (the EDF statistics are upper tail tests), none of the four distributions can provide a satisfactory modelling of

the unconditional transaction-aggregated durations. However, the generally decreasing EDF statistics (especially for Gamma distributions) imply that there is some aggregational characteristics of the transaction-aggregated durations. Given the large sample of tick-by-tick durations, we can further investigate this aggregational characteristic by increasing T that is the fixed number used to aggregate the tick-by-tick durations (the trading day with least tick-by-tick durations still have 1214888 observations).

		Exponential			Weibull			Gamma			Generalized Pareto		
	Sample size	D	W^2	A^2	D	W^2	A^2	D	W^2	A^2	D	W^2	A^2
T=13	6596215	0.311	2211.59	Inf	0.0812	12430	Inf	0.088	7669	Inf	0.168	57340	4E+05
T=67	1199206	0.048	1305.789	Inf	0.0367	514	Inf	0.039	498.2	Inf	0.036	500.2	Inf
T=134	595032	0.037	229.8506	Inf	0.0316	185.4	Inf	0.026	148.4	Inf	0.020	60.83	Inf
T=400	198261	0.105	630.4093	Inf	0.0421	127.9	Inf	0.033	73.5	Inf	0.0167	19.64	Inf
T=800	98942	0.131	504.7833	3183.103	0.0435	72.28	Inf	0.035	39.1	Inf	0.021	14.88	98.15
T=1200	65893	0.143	404.4711	2496.856	0.0446	51.16	Inf	0.036	26.7	Inf	0.022	11.24	73.21
T=2400	32876	0.159	254.6905	1537.22	0.0471	26.82	Inf	0.037	13.5	Inf	0.025	6.788	44.59

Table 3.5 The fitting of Exponential, Weibull, Gamma and generalized Pareto distribution to transaction-aggregated durations

		Exponential			Weibull			Gamma			Generalized Pareto		
	Sample size	D	W^2	A^2	D	W^2	A^2	D	W^2	A^2	D	W^2	A^2
T=4000	19663	0.171	175.61	1043.05	0.048	17.07	110.84	0.037	8.416	46.632	0.028	4.693	30.58
		[0.014]	[0.586]	[2.862]	[0.009]	[0.439]	[2.421]	[0.008]	[0.439]	[2.258]	[0.011]	[0.571]	[3.033]
T=8000	9763	0.187	103	600.063	0.04991	8.855	57.609	0.041	4.256	23.748	0.033	2.76	17.73
		[0.014]	[0.496]	[2.704]	[0.014]	[0.401]	[2.195]	[0.012]	[0.409]	[2.249]	[0.0148]	[0.593]	[3.032]
T=12000	6468	0.194	74.206	427.937	0.05004	6.007	38.879	0.039	2.822	15.851	0.033	1.927	12.59
		[0.016]	[0.421]	[2.371]	[0.015]	[0.391]	[2.091]	[0.015]	[0.438]	[2.237]	[0.015]	[0.348]	[2.061]
T=24000	3162	0.205	41.168	233.888	0.05101	2.634	17.323	0.038	1.154	6.5967	0.035	1.025	6.639
		[0.024]	[0.461]	[2.449]	[0.023]	[0.482]	[2.347]	[0.022]	[0.431]	[2.355]	[0.022]	[0.353]	[2.001]
T=40000	1854	0.217	26.504	149.038	0.05693	1.599	10.046	0.04	0.734	4.0912	0.04	0.782	4.904
		[0.032]	[0.507]	[2.535]	[0.032]	[0.539]	[2.791]	[0.033]	[0.631]	[3.423]	[0.031]	[0.401]	[2.129]
T=60000	1193	0.226	19.095	106.202	0.04383	0.807	5.3345	0.035	0.354	2.2714	0.039	0.497	3.409
		[0.041]	[0.625]	[3.192]	[0.035]	[0.352]	[2.003]	[0.041]	[0.555]	[2.849]	[0.037]	[0.371]	[2.105]

Table 3.6 The fitting of Exponential, Weibull, Gamma and generalized Pareto distribution to transaction-aggregated durations

Table 3.6 summarizes the goodness of fit for the $T-4000$, $T-8000$, $T-12000$, $T-24000$, $T-40000$ and $T-60000$ transaction-aggregated durations. The critical values are reported in the brackets and the significance level is 5%. The EDF statistics are very sensitive even when the sample is relatively small. The details of the power studies with respect to EDF statistics can be found in the classic paper of Stephens (1974). There are two interesting findings that can be concluded from Table 3.6. First, compared with other distributions under consideration, the Gamma distributions fit the transaction-aggregated duration much better in the sense that the corresponding EDF statistics are closer to the critical values.

Second, the distributions of transaction-aggregated durations approach the Gamma distribution as the number used to aggregate to tick-by-tick durations, T , increases. For instance, for $T-40000$ transaction-aggregated durations, the EDF statistics are 0.04, 0.734 and 4.0912 for D , W^2 and A respectively that are very close to the corresponding critical values of 0.033, 0.631 and 3.423. Meanwhile, for $T-60000$ transaction-aggregated durations, the three EDF statistics are 0.035, 0.354 and 2.2714. The corresponding critical values are 0.041, 0.555 and 2.849 respectively. Thus, the null hypothesis of a Gamma distributed sample cannot be rejected at the significance level of 5%. The following Figure 3.17 and Figure 3.18 present respectively the empirical cumulative distribution functions of $T-24000$ and $T-40000$ durations together with corresponding estimated CDF of Gamma distributions.




Figure 3.17 The empirical cumulative distribution functions of T-24000 transaction-aggregated durations




Figure 3.18 The empirical cumulative distribution functions of T-40000 transaction-aggregated durations

The above empirical results suggest the existence of the aggregational characteristics of tick-by-tick durations. It is well-known that the intraday financial transactions are not only subject to “time-of-day” effect but also “day-of-week” effect. Therefore, I further investigate the aggregational characteristics for tick-by-tick durations within each trading day. The procedures are similar to the one that I apply for studying aggregational normality of intraday returns in Chapter 1. The purpose is to explore whether the intraday durations exhibit aggregational characteristics.

Specifically, I carry out the test for the null hypothesis of Gamma distribution using EDF statistics D , W^2 and A^2 on each day’s transaction-aggregated durations. Transaction-aggregated durations within a trading day are considered to follow Gamma distributions if all the three statistics D , W^2 and A^2 cannot reject the null hypothesis at significance level of 5%. As suggested by the simple size in Table 3.7, the choices of T generate enough data for intraday durations.

	Average daily sample size	Exponential		Weibull		Gamma	
		N_0	N_0/N	N_0	N_0/N	N_0	N_0/N
T=13	26175	0	0%	0	0%	0	0%
T=67	4758	0	0%	0	0%	0	0%
T=134	2361	2	0.79%	21	8.33%	16	6.35%
T=400	786	0	0%	48	19.04%	133	52.78%
T=800	392	0	0%	81	32.14%	162	64.29%
T=1200	261	0	0%	112	44.44%	175	69.44%
T=2400	130	0	0%	160	63.49%	199	78.97%

Table 3.7 The fitting of Exponential, Weibull and Gamma distributions to within-day transaction-aggregated duration

The N_0 in Table 3.7 represents the number of days for which the EDF statistics cannot reject the null hypothesis that the within-day transaction-aggregated durations come from the family of corresponding distribution and N is the total number of trading days in our sample. Table 3.7 have several important implications.

First, the tick-by-tick duration presents aggregational characteristics regarding its empirical unconditional distributions. Although it cannot be described effectively by the distributions commonly used in duration modellings such as exponential, Weibull and Gamma distributions, when aggregated by tick-by-tick durations, the distribution of transaction-aggregated durations approaches a shape similar to Gamma and Weibull distributions.

Second, compared with exponential distributions, Weibull and Gamma distributions seem to be more suitable for modelling transaction-aggregated durations. Table 3.7 gives very clear evidence against the use of exponential distributions in modelling transaction-aggregated durations. In contrast, the Gamma distribution seems to be promising in the modelling of the unconditional transaction-aggregated durations. For instance, there are 162 trading days (64.29% of the whole sample of 252 trading days) for which the EDF statistics cannot reject the null hypothesis that the T -800 (each T -800 duration is aggregated from 800 tick-by-tick durations) durations within the trading day follow Gamma distributions. In other words, the probability of the T -800 durations within a trading day to follow a Gamma distribution is 62.29%. With respect to the T -2400 transaction-aggregated durations within a trading day, the probability of them to follow a Gamma distribution is 78.97%.

Finally, as reported previously, since tick-by-tick durations present characteristics that are directly determined by market operational details, when measured at the highest frequency (T -13, T -67 and T -134 durations), the transaction-aggregated durations cannot be effectively described by the continuous distributions that are commonly applied in duration modellings.

In conclusion, in this section I provide an investigation with respect to the unconditional

distributional properties with respect to tick-by-tick durations. I show that the pattern of tick-by-tick durations of SPY are largely influenced by market microstructural issues discussed in section 3.3.1. Besides, I identify that the tick-by-tick durations present aggregational characteristics and Gamma distributions seem to serve as a good candidate for modelling unconditional transaction-aggregated durations.

3.4 ACD modelling of durations

In this section I present the empirical results with respect to the ACD duration modelling. I use the transaction-aggregated durations instead of the tick-by-tick durations since that, as discussed in Section 3.3.4, the SPY tick-by-tick durations are largely influenced by market microstructural issues. Unfortunately, I cannot explicitly exclude the effect of such issues using my data. Thus, a direct use of ACD models on the tick-by-tick duration of SPY might provide misleading results.

3.4.1 Data preparation: Intraday seasonality and re-initialization

Studies regarding intraday financial data have widely documented the strong intraday pattern over the trading day (Jain and John 1988, Andersen and Bollerslev 1997 and Bollerslev and Domowitz 1993). Before applying ACD models to describe the transaction-aggregated duration dynamics, it is necessary to exclude this intraday periodic pattern from the raw transaction-aggregated durations. Since the transaction-aggregated durations are calculated by the sum of tick-by-tick durations, excluding the intraday seasonality from the raw transaction-aggregated duration is equivalent to excluding the intraday seasonality from the raw tick-by-tick durations. In other words, I can first de-seasonalize the raw tick-by-tick durations and further build the corresponding transaction-aggregated durations by taking the sum of seasonal adjusted tick-by-tick durations.

In Section 3.3.3, I investigate the intraday pattern of tick-by-tick durations of SPY and briefly introduce the method that I use to exclude the “time-of-day” effect from the raw tick-by-tick intraday durations. I now elaborate the procedures.

I first remove all tick-by-tick durations that exceed 3s or equal to zero from the sample (the 0.99 quantile of the tick-by-tick sample without zeros is 2.226 second). Recall the assumption of (3.49),

$$w_i = \tilde{w}_i s(t_{i-1}),$$

where $s(t_i)$ is the intraday periodicity component at t_i and $w_i = t_i - t_{i-1}$ is the i th tick-by-tick duration. The intraday periodic component $s(t_i)$ is commonly modelled by cubic splines with nodes set on some fixed interval over the trading day. The constant on each node is given by the observed sample mean over the corresponding time interval. A clear example is presented by Figure 3.9 in Section 3.3.3.

However, as one might have noticed, the intraday tick-by-tick duration pattern in Figure 3.9 is very different from the one in Figure 3.3. This contradiction is mainly caused by the extremely high liquidity of SPY. Generally, the nodes of the cubic splines are setting on relatively long intervals over the trading day to make sure the spline is smooth. For instance, in Engle and Russell (1998), the nodes are setting on each hour of the trading day. Bauwens and Giot (2000) applied thirty minutes intervals. For tick-by-tick durations of SPY without zeros, the 0.99 quantile of the sample is only 2.226 second (the 0.99 quantile of the raw tick-by-tick sample with zeros is only 1.441s). Consequently, if a long interval is used, the long tick-by-tick durations that usually occur after middle day might bias the sample mean to a large extent. This is exactly the case of Figure 3.9 that is plotted using whole tick-by-tick sample.

I further notice that there is a dilemma when using ACD models to describe tick-by-tick duration dynamics of stocks with high liquidity like SPY. From economic point of view, the long durations are generally believed to convey important information as discussed in section 3.2.2.3. Nevertheless, for a symbol like SPY, most of the transactions are liquidity transactions that refer to the buy or sell for liquidity and therefore most of the tick-by-tick durations are of extremely small values. Consequently, there will be some “jumps” in the duration series. Thus, if one wants to retain as many observations as possible in the sample for ACD modellings of duration dynamics, the ACD specifications are challenged to

capture these “jumps”.

Besides, the characteristic of tick-by-tick durations for stocks with extremely high liquidity also suggest that the commonly used de-seasonal technique of cubic splines might need a careful review. In addition to the technique difficulty of choosing proper time intervals, a more serious issue is that the averages over time intervals of trading days might not be periodic. For instance, as discussed in Section 3.3.3, the averages of tick-by-tick durations over 1-minute intervals of different trading days seem to have little periodic pattern in the correlogram.

Although I believe that studies in the above two directions can yield valuable knowledge on the modelling of duration dynamics, it unfortunately is beyond the scope of this chapter. In this chapter, I simply rule out the durations that exceeds 3s as mentioned previously so that I can focus on the comparison of different distribution assumptions.




Figure 3.19 Cubic spline function with 15-minute nodes of SPY tick-by-tick durations

Figure 3.19 presents the cubic spline function with 15-minute nodes. The constant on each node is given by the observed sample mean of tick-by-tick durations over the corresponding 15-minute interval. Now the spline line is much more in line with the

intraday tick-by-tick pattern in Figure 3.3. The market opening is very active as expected with transactions occurring, on average, every 0.1s. In contrast, the middle of day, around 13:00, has the longest averaged tick-by-tick durations around 0.27s. The trading intensity rises again until the closing of the market. Specifically, near 16:00, the transactions occur again almost every 0.08s.

I now exhibit the descriptive statistics of the raw transaction-aggregated durations and the seasonal filtered transaction-aggregated durations respectively.

Raw durations	Sample size	Mean	Standard Deviation	q (0.05)	q (0.25)	q (0.5)	q (0.75)	q (0.95)	Ljung-Box statistics
Tick durations	30542307	0.162	0.363	0.001	0.001	0.011	0.129	0.897	7891637
T-13 durations	2349290	2.104	2.162	0.069	0.533	1.391	2.977	6.567	4607113
T-67 durations	455735	10.847	7.910	1.675	4.772	8.898	15.136	26.474	2456680
T-134 durations	227801	21.699	14.405	4.150	10.443	18.482	30.140	49.768	1558918
T-400 durations	76233	64.831	38.589	14.885	34.290	57.538	89.434	137.979	629368
T-800 durations	38054	129.839	73.257	32.307	71.638	117.021	177.719	266.746	295839
T-1200 durations	25327	195.026	106.965	50.971	109.960	177.542	266.873	394.675	169913
T-2400 durations	12608	391.299	205.066	110.154	228.128	361.458	531.184	769.135	89904

Table 3.8 Descriptive statistics of the raw transaction-aggregated durations

Filtered durations	Sample size	Mean	Standard Deviation	q (0.05)	q (0.25)	q (0.5)	q (0.75)	q (0.95)	Ljung-Box statistics
Tick durations	30542307	0.817	1.790	0.004	0.008	0.059	0.697	4.387	6053335
T-13 durations	2349290	10.620	10.265	0.373	2.955	7.524	15.218	31.411	3350750
T-67 durations	455735	54.744	36.298	9.633	26.622	47.395	75.783	124.091	1808628
T-134 durations	227801	109.511	65.120	24.356	58.369	98.637	149.743	230.738	1142505
T-400 durations	76233	327.132	170.805	89.073	191.315	307.626	440.641	634.629	479633
T-800 durations	38054	654.958	320.347	194.702	400.873	626.931	873.049	1217.990	258894
T-1200 durations	25327	983.445	464.559	306.379	615.816	947.819	1303.624	1796.435	169913
T-2400 durations	12608	1971.052	881.642	650.639	1277.236	1927.322	2589.364	3490.888	55324

Table 3.9 Descriptive statistics of the filtered transaction-aggregated durations

As suggested by the Ljung-Box statistics in Table 3.9, the seasonal filtered transaction-aggregated durations still present strong autocorrelations. It suggests the periodic pattern of Figure 3.18 is not the only reason behind the strong autocorrelations of the raw transaction-aggregated durations. In general, compared with Table 3.8, the LB statistics in Table 3.9 decrease as expected, which suggests the cubic spline function indeed removes some trend pattern from the raw durations.

Another important issues that I address here is the re-initialization of tick-by-tick durations in each trading day. Clearly, durations are calculated consecutively from day to day. For instance, the first tick-by-tick duration of a specific trading day will always be the time difference between the first and second transactions in that trading day. Hence, I re-initialize for tick-by-tick durations within each trading day. Specifically, when applying ACD models, the conditional expectation of the first transaction-aggregated durations of each trading day will be the average of the transaction-aggregated durations over the first 15 minutes of the trading day. Consequently, the ACD process of each trading day starts at 9:45 that is in line with the cubic spline function in Figure 3.19.

3.4.2 ACD modelling of transaction-aggregated duration dynamics

In this section I present the empirical results of ACD modelling on transaction-aggregated durations. With respect to the conditional expectation mean specifications, I realize that there is a wide range of possibilities as discussed in Section 3.2. However, due to the computational difficulties arising with the complexity of those ACD models, I focus on the comparison of different distributional assumptions based on our prior findings regarding the unconditional distributional properties of tick-by-tick durations.

Thus, I choose the standard ACD (m, q) of (3.13),

$$\psi_i = \omega + \sum_{j=1}^m \alpha_j w_{i-j} + \sum_{j=1}^q \beta_j \psi_{i-j}.$$

I select the following two specifications with different orders of lags: ACD (1,1) and ACD

(2,1). Regarding the distributional assumptions of the errors, the three distributions that I consider are the exponential distribution, Weibull distribution and Gamma distributions. Thus, there are six ACD specifications in total.

The diagnostics are: Log-likelihood (LL), Bayes Information Criterion (BIC) and Ljung-Box statistics (LB) with respect to 20 lags of model residuals. LL and BIC are applied to compare the different specifications and meanwhile LB is used to test the autocorrelations of residuals.

I use the T -13, T -67 T -134 and T -400 transaction-aggregated durations to fit the selected ACD specifications. Each transaction-aggregated duration series is normalized by sample means before the estimation. Given the extremely high value of LB statistics in Tables 3.8 and 3.9, it is interesting to see whether the ACD models can meet this challenge.

T -13	Exponential		Weibull		Gamma	
	EACD (1,1)	EACD (2,1)	WACD (1,1)	WACD (2,1)	GACD (1,1)	GACD (2,1)
ω	0.0013 (0.0002)	0.0036 (0.0002)	0.0093 (0.0002)	0.0029 (0.0001)	0.0014 (0.0002)	0.0091 (0.0004)
α_1	0.0183 (0.0012)	-0.0013 (0.0006)	0.0557 (0.0011)	0.0531 (0.0009)	0.0132 (0.0008)	0.0172 (0.0007)
β_1	0.9445 (0.0007)	0.9501 (0.0005)	0.9379 (0.0011)	0.9212 (0.0012)	0.9762 (0.0003)	0.9622 (0.0003)
α_2		0.0087 (0.0011)		0.0427 (0.0011)		0.0105 (0.0013)
Diagnostics						
Sample size	2349290	2349290	2349290	2349290	2349290	2349290
LL	-2011007	-1993041	-1966524	-1947097	-1773042	-1723578
BIC	4022058	3986126	3933092	3894238	3546128	3447221
LB	51102	50502	48324	47032	43402	42456

Table 3.10 ACD estimation of T-13 transaction-aggregated durations

T-67	Exponential		Weibull		Gamma	
	EACD (1,1)	EACD (2,1)	WACD (1,1)	WACD (2,2)	GACD (1,1)	GACD (2,1)
ω	0.0319 (0.0008)	0.0075 (0.0002)	0.0186 (0.0008)	0.0056 (0.0003)	0.0173 (0.0004)	0.0063 (0.0002)
α_1	0.2094 (0.0031)	0.1214 (0.0018)	0.1497 (0.0031)	0.1112 (0.0022)	0.2032 (0.0026)	0.1672 (0.0026)
β_1	0.7593 (0.0036)	0.8765 (0.0024)	0.8320 (0.0026)	0.8718 (0.0031)	0.7762 (0.0035)	0.7822 (0.0034)
α_2		-0.0088 (0.0030)		-0.0928 (0.0038)		0.0345 (0.0009)
Diagnostics						
Sample size	455735	455735	455735	455735	455735	455735
LL	-398949	-397742	-239314	-233694	-167542	-155358
BIC	797937	795536	478667	467440	335123	310768
LB	10773	9856	8883	8241	7734	7413

Table 3.11 ACD estimation of T-67 transaction-aggregated durations

T-134	Exponential		Weibull		Gamma	
	EACD (1,1)	EACD (2,1)	WACD (1,1)	WACD (2,1)	GACD (1,1)	GACD (2,1)
ω	0.0203 (0.0007)	0.0045 (0.0003)	0.0303 (0.0019)	0.0035 (0.0005)	0.0224 (0.0008)	0.009 (0.0005)
α_1	0.3767 (0.0038)	0.4655 (0.0026)	0.3029 (0.0058)	0.4084 (0.0034)	0.2152 (0.0019)	0.1856 (0.0012)
β_1	0.6078 (0.0042)	0.8341 (0.0045)	0.6641 (0.0069)	0.9103 (0.0076)	0.7362 (0.0046)	0.8122 (0.0049)
α_2		-0.3026 (0.0049)		-0.3226 (0.0086)		-0.0305 (0.0068)
Diagnostics						
Sample size	227801	227801	227801	227801	227801	227801
LL	-201301	-200758	-90405	-88477	-67304	-64311
BIC	402639	401565	180847	177003	134645	128671
LB	5728	5523	3628	3356	2785	2556

Table 3.12 ACD estimation of T-134 transaction-aggregated durations

T-400	Exponential		Weibull		Gamma	
	EACD (1,1)	EACD (2,1)	WACD (1,1)	WACD (2,1)	GACD (1,1)	GACD (2,1)
ω	0.0266 (0.0014)	0.0030 (0.0003)	0.0210 (0.0035)	0.0018 (0.0003)	0.0286 (0.0015)	0.0173 (0.0028)
α_1	0.4500 (0.0081)	0.5393 (0.0043)	0.2187 (0.0205)	0.4299 (0.0064)	0.2232 (0.0224)	0.2172 (0.0194)
β_1	0.5275 (0.0088)	0.8624 (0.0066)	0.7565 (0.0242)	0.9390 (0.0041)	0.7282 (0.0271)	0.8522 (0.0069)
α_2		-0.4038 (0.0076)		-0.3715 (0.0070)		-0.1105 (0.0048)
Diagnostics						
Sample size	76233	76233	76233	76233	76233	76233
LL	-68454	-68259	-20289	-16527	-12733	-11457
BIC	136941	136562	40611	33098	25499	22958
LB	2382	2134	1678	1622	1328	1246

Table 3.13 ACD estimation of T-400 transaction-aggregated durations

Tables 3.10 – 3.13 summarize the ACD estimation of the four transaction-aggregated durations. The estimation is based on MLE. Standard errors are calculated by the method of White and are given in parentheses. I do not impose any nonnegative constraints on the parameters in the estimation. For this exact reason, there are negative parameters for α_2 when the ACD (2,1) model is estimated. Nevertheless, in these cases, they are compensated by rising value of α_1 and we do not observe any negative ψ_i in the estimation. It is very intriguing since such α_2 's is of non-trivial values. I now summaries the findings from the four tables.

With respect to the specifications, I first notice that ACD (2,1) models outperform the ACD (1,1) models dominantly. The BIC and LL of ACD (2,1) models are constantly better than ACD (1,1) models (higher LL and lower BIC). It implies that the ACD (1,1) models might be over simplistic for modelling SPY durations. Further evidences can be found in the LB statistics. Compared with LB statistics of the seasonal filtered transaction-aggregated

durations in Table 3.9, the LB statistics of residuals of ACD models are reduced massively. Nevertheless, the LB statistics still clearly exceed the critical values of 31.41. In other words, the residuals still exhibit significant autocorrelations. This might not be very surprising since, as I mentioned previously, the SPY seasonal filtered transaction-aggregated durations exhibit very strong autocorrelations (Table 3.9) that imply the existence of a trend. In this case, an ARMA-GARCH specification for the logarithm of durations might provide better description of the autocorrelations. Consequently, higher orders of the lags in ACD models are preferred to give more explanatory powers over the duration dynamics. For the compactness of the Tables, I do not include the T statistics into the table (the T statistics of w_2 in the ACD (2,1) exceed 2 for most of the cases).

Second, the parameter α in general are of very small values and the β are close to unit. For instance, for the T-13 durations, the α_1 of EACD (1,1), WACD (1,1) and GACD (1,1) are 0.0013, 0.0093 and 0.0014 respectively and the β of the three models are 0.9445, 0.9379 and 0.9762. It suggests that the persistence of the transaction-aggregated duration process is very strong (which is in line with the previous empirical findings in Section 3.3.2). Also, in general the α gradually increase and β decreases with the frequencies of transaction-aggregated durations. Consequently, recent observations now contribute more to the conditional expectations ψ_i .

Finally, based on the LL and BIC statistics, the ACD models with Gamma innovations outperform the ACD models with other two innovations very obviously. For each series of transaction-aggregated durations, the GACD models always have the highest LL and the lowest BIC, suggesting that they fit the data best. It further confirms the previous empirical finding with respect to the unconditional distributional properties of transaction-aggregated durations.

The following Figures 3.20 and 3.21 exhibit the correlogram and histogram of the residuals of GACD (1,1) models for T -400 durations. Figure 3.20 clearly suggests that the residuals present significant autocorrelations at the first several lags. It should be mentioned that I select only the simplest linear ACD models. The possible nonlinear dependence cannot be

captured by the linear ACD (m,q) model. Besides, adding explanatory variables into the conditional expectation specification might be able to improve the fit as well.




Figure 3.20 The correlogram of T-400 GACD (1,1) residuals




Figure 3.21The histogram of T-400 GACD (1,1) residuals

3.5 Conclusions

In this chapter I study the modelling of durations between transactions. I start by discussing and analyzing both theoretical and empirical works on the ACD modelling of duration dynamics. I further conduct a data analysis with respect to the tick-by-tick durations of SPY. Based on the empirical evidence, we show that the tick-by-tick transactions for stocks with ultra-high liquidity like SPY are largely affected by market microstructural issues such as the market's second-by-second operational details. Thus, studies using such data should be very careful when trying to invoke a parametric specification for modelling duration dynamics. Besides, I showed that the tick-by-tick durations have aggregational characteristics with respect to its unconditional distribution. Specifically, the empirical distributions of sums of tick-by-tick durations approximate a shape similar to the Gamma family. Standard ACD models with different distribution assumptions are applied to model the transaction-aggregated durations. The empirical results suggest that the standard ACD models cannot fully capture the interdependence structure of SPY durations. Nevertheless, the ACD specifications with Gamma distributed innovations outperforms the ACD specifications with other innovations significantly. Thus, Gamma distributions may serve as a good starting point in duration modellings. Besides, the future research direction could be the “day of week” seasonality in durations.

Appendix A.

The classic paper of Stephens in 1974 presents a detailed discussion with respect to the empirical distribution function statistics such as the statistics D (derived from D^+ and D^-), W^2 , V , U^2 and A^2 . Given some random sample $x_1, x_2, x_3, \dots, x_n$ and the null hypothesis H_0 that the random sample comes from a distribution with distribution function $F(x)$, the EDF statistics compare the $F(x)$ with the empirical distribution function $F_n(x)$. In our case, we estimate the u and σ^2 by the average of the sample \hat{x} and $\hat{\sigma}^2 = \sum_i (x_i - \hat{x}) / (n - 1)$.

The sample are supposed to be indexed by values in ascending order, $x_1 \leq x_2 \leq x_3 \leq \dots \leq x_n$. The test procedures are organized as follows.

- (a) Estimate parameters u and σ^2 as discussed previously.
- (b) Calculate z_i 's as $z_i = F(x_i)$, where $F(x)$ is the distribution function of normal distribution with u and σ^2 estimated in (a).
- (c) Calculate D , W^2 , V , U^2 and A^2 as follows:

- 1. The Kolmogorov statistics D , D^+ and D^- :

$$D^+ = \max_{1 \leq i \leq n} \left[\left(\frac{i}{n} \right) - z_i \right];$$

$$D^- = \max_{1 \leq i \leq n} \left[- \left(\frac{i-1}{n} \right) + z_i \right];$$

$$D = \max (D^+, D^-).$$

- 2. The Cramér-von Mises statistics W^2 :

$$W^2 = \sum_{i=1}^n \left[z_i - \frac{(2i-1)}{2n} \right]^2 + 1/12n.$$

- 3. The Kuiper statistic V :

$$V = D^+ + D^-.$$

- 4. The Waston statistics U^2 :

$$U^2 = W^2 - n \left(\sum_{i=1}^n \frac{z_i}{n} - \frac{1}{2} \right)^2.$$

5. The Andersen-Darling statistics A^2 :

$$A^2 = -\frac{\{\sum_{i=1}^n (2i-1)[\ln z_i + \ln(1-z_{n+1-i})]\}}{n} - n.$$

I produce my table of critical values by Monte-Carol simulation. The null hypothesis H_0 that the sample comes from a normal distribution is rejected if the values of the statistics are larger than the critical values given corresponding significance levels. The test is the traditional upper-tail test.

Statistics	Significance level		
	5%	2.5%	1%
D	0.895	0.955	1.035
V	1.489	1.585	1.693
W^2	0.126	0.148	0.178
U^2	0.116	0.136	0.163
A^2	0.787	0.918	1.092

Table 1.10(a) Critical values of EDF statistics for the null hypothesis of normal distribution.

Appendix B.

The standard Central Limit Theory is the cornerstone of probability theory and statistics. It asserts that:

If $\{X_t\}$ be a sequence of identically distributed independent random variables with $E(X_t) = u$ and $Var(X_t) = \sigma^2$ for all t . Then,

$\lim_{t \rightarrow \infty} P\left(\frac{\sum_{t=1}^n X_t - nu}{\sqrt{n\sigma}} \leq x\right) = \Phi(x)$, where Φ is the cumulative distribution function of the standard normal distribution.

In many economic circumstances, the *i.i.d* assumption cannot be assumed to hold easily. Thus, an issue of both practical and theoretical importance is that:

Assume that $\{X_t\}_{t \in \mathbb{Z}}$ is a process with zero mean $E(X_t) = 0$ and finite second moment $E(X_t^2) < \infty$. Define the partial sum $S_n = \sum_{t=1}^n X_t$ and its normalized variance $s_n^2 = \frac{Var(S_n)}{n}$. Whether or not we can apply the Central Limit Theorem to have $\lim_{n \rightarrow \infty} P(S_n \leq x\sqrt{ns_n^2}) = \Phi(x)$ where Φ is the cumulative distribution function of the standard normal distribution. Moreover, if the central limit theorem holds, what is the possible convergence rate. Most of the literatures regarding the convergence rate use the Kolmogorov metric as the metric of underlying $\{X_t\}$. Define $\Delta_n(x) = |P(S_n \leq x\sqrt{ns_n^2}) - \Phi(x)|$ and further $\Delta_n = \sup \{\Delta_n(x) | x \in \mathbb{R}\}$. Hörmann (2009) provides the bounds for normal approximation error Δ_n for dependent $\{X_t\}$. An important class of dependent $\{X_t\}$ is the ARCH/GARCH process.

In the last decades, the ARCH/GARCH model are widely applied for the modelling of time-varying volatility. In 1997, Duan introduces a general functional form of GARCH model, the so-called augmented GARCH (1,1) process, that contains many existing GARCH models as special cases. The augmented GARCH (1,1) process is given by:

$$y_t = \sigma_t \varepsilon_t, \\ \Lambda(\sigma_t^2) = c(\varepsilon_{t-1})\Lambda(\sigma_{t-1}^2) + g(\varepsilon_{t-1}),$$

where Λ , c and g are real-valued measurable functions and $\{y_t\}_{t \in \mathbb{Z}}$ is a random variable and $\{\varepsilon_t\}_{t \in \mathbb{Z}}$ is an *i.i.d* sequence. In order to solve σ_t^2 , the definition requires the existence of Λ^{-1} . Aue et al. (2006) discuss the necessary and sufficient conditions for the augmented GARCH (1,1) model to have a strictly stationary and non-negative solution for σ_t^2 . For the special case of GARCH (1,1),

$$\sigma_t^2 = \omega + \beta \sigma_{t-1}^2 + \alpha y_{t-1}^2 = \omega + [\beta + \alpha \varepsilon_{t-1}^2] \sigma_{t-1}^2 .$$

It requires that $E(\log|\beta + \alpha \varepsilon_0^2|) < 0$. Hörmann (2008) analyzes the dependence structure and asymptotical properties of the augmented GARCH process and shows that m -dependent approximations $\{Y_{tm}\}$ to the original sequence of $\{Y_t\}$ can be obtained. Specifically, $\|Y_{tm} - Y_t\|_2 < \text{const} \cdot \rho^m$ ($\rho < 1$), where $\|\cdot\|_2$ is the L^2 norm. He then uses the m -dependent approximations to construct the proof of convergence to normal distribution and corresponding convergence rate. Details of the proof can be found in Hörmann (2008).

Based on this result, consider the ARMA (p, q)-GARCH (1,1) model:

$$\begin{aligned} x_t &= \phi_1 x_{t-1} + \dots + \phi_p x_{t-p} + \theta_1 y_{t-1} + \dots + \theta_q y_{t-q}, \\ y_t &= \sigma_t \varepsilon_t, \\ \sigma_t^2 &= \omega + \beta \sigma_{t-1}^2 + \alpha y_{t-1}^2. \end{aligned}$$

If a strictly stationary and casual solution of the equation $x_t = \phi_1 x_{t-1} + \dots + \phi_p x_{t-p} + \theta_1 y_{t-1} + \dots + \theta_q y_{t-q}$ exists, Brockwell and Davis (1991) show that the solution can be presented as a linear process $\sum_i \psi_i y_{t-i}$ where the coefficients ψ_i exponentially decay as the lags. Then the convergence of the distribution of normalized sums of x_t 's to the normal distribution can be verified. According to Hörmann (2009), if $p \in (2, 3]$ moments exist for x_t , $\Delta_n = O((\log n)^{p-1} n^{1-\frac{p}{2}})$.

However, our tick-by-tick return data presents many erratic statistical features that cannot be explained by the ARMA-GARCH models. The fit of ARMA-GARCH models using tick-by-tick returns gives poor results. The required stationarity is violated since the fit

gives the value of $\alpha + \beta$ that exceeds 1. More importantly, the raw tick-by-tick returns are seriously affected by the market microstructural issues such as “split-transactions” which we cannot analyze since the bid and ask prices are not included in the dataset. Thus, a theoretical modelling of the tick-by-tick price dynamics seem to be impractical.

Appendix C

The following Matlab program “tickstandardize” estimates the periodicity pattern for given trading interval. In practice, it is used with other programs jointly. For instance, if I want to filter the transaction-aggregated returns T-400, I will use the T-400 return series and corresponding tick-by-tick table in a for-loop to find out its periodicity pattern. The for-loop will search the whole tick-table. For time-aggregated return series, the computation is much simpler since the for-loop will be executed for fixed time interval.

```
function [r1] =tickstandardize(Ticktable,tradetimes,r)

% tradetimes and r are one day's data for a specific day, Ticktable is
% 252*4 cell table for tick transactions

n=length(r); % number of returns in r series

for i=1:n
    s2=tradetimes(i+1);
    s1=tradetimes(i); % trade interval is [s1 s2] for return r(i)
    v=[];

    for j=1:252

        tickprices=Ticktable{1}{j};
        ticktimes=Ticktable{2}{j};

        transactions=[]; % This variable is for the transactions that occurs between s1 and s2 on
        day j

        for k=1:length(ticktimes)

            if ticktimes(k)>=s1 && ticktimes(k)<=s2
                transactions=[transactions;tickprices(k)];
            end
        end
    end
end
```

```

elseif ticktimes(k)>s2
break
end
end
standardizer=log(transactions(2:end))-log(transactions(1:end-1));
v(j)=sum(standardizer.^2); % use realized variance to estimate the periodicity
end
sv=(sum(v)/252)^(1/2);
r1(i)=r(i)/sv;
end
end

```

The following Matlab program “tfiteration” uses the conditional daily variance and intraday periodicity to normalize the intraday transaction-aggregated returns.

```
function [p,t,r,adjustr,ar] = tfiteration(data,tickt,tickp,T,v)
% This function uses the conditional daily variance and intraday
% periodicity to normalize the returns.
% data is the spyfiltered data,tickt,tickp are the corresppongding
% ticktable, T is the frequency index

dayindex=datasort(data{2});

for i=1:length(dayindex)-1

dayprices=data{4}(dayindex(i)+1:dayindex(i+1));
daytradetimes=data{3}(dayindex(i)+1:dayindex(i+1)); % extract the day data

[p{i},t{i},r{i},~]=tsprices(dayprices,daytradetimes,T); % get the intraday transaction-
aggregated returns.
s=[];
for j=1:length(r{i})
t1=t{i}(j);
t2=t{i}(j+1); % [t1,t2] is the timeinterval of jth intraday return
prices1=pricelock(t1,tickp,tickt);
prices2=pricelock(t2,tickp,tickt);
ticksquare=log(prices2./prices1).^2;
s(j)=mean(ticksquare./v); %s(j) is actually s^2(j)
end
```

```

adjustr{i}=(r{i}/sqrt(v(i)))./sqrt(s);
end
ar=[];
for i=1:length(dayindex)-1
ar=[ar, adjustr{i}];
end
end

```

This Matlab program “pricelock” searches the tick-by-tick transaction records for transactions occurring in specific time interval.

```
function [prices] = pricelock(timing,tickp,tickt)
```

```
for i=1:length(tickp)
```

```
    t=tickt{i}; %get the ith day tick time table
```

```
    n=length(t); %n is the length of transactions
```

```
    right=t(t>=timing);
```

```
    left=t(t<=timing);
```

```
    if isempty(right)
```

```
        pright=tickp{i}(end);
```

```
    else
```

```
        pright=tickp{i}(n-length(right)+1);
```

```
    end
```

```
    if isempty(left)
```

```
        pleft=tickp{i}(1);
```

```
    else
```

```
        pleft=tickp{i}(length(left));
```

```
    end
```

```
    prices(i)=(pright+pleft)/2;
```

```
end
```


Appendix D

This appendix includes the introduction of the MLE estimation for the exponential distribution, the Weibull distribution, the Gamma distribution and the generalized Pareto distribution. The optimization of the loglikelihood functions is conducted by Matlab solver based on iteration. It also includes the Matlab code for Monte-Carlo simulations that is used to produce the table of critical values for EDF statistics. The procedure is similar to the one in Appendix A.

Let $\{x_1, \dots, x_n\}$ be the random sample from some distribution. The exponential distribution has the probability density function

$$f(x; \lambda) = \begin{cases} \lambda e^{-\lambda x} & x \geq 0, \\ 0 & x < 0. \end{cases}$$

The loglikelihood function therefore is $l(\lambda; x_0, x_1, \dots, x_n) = n \ln(\lambda) - \lambda \sum_{i=1}^n x_i$.

With respect to the Weibull distribution, it has the probability density function

$$f(x; \lambda, k) = \begin{cases} \frac{k}{\lambda} \left(\frac{x}{\lambda}\right)^{k-1} \exp\left(-\left(\frac{x}{\lambda}\right)^k\right) & x \geq 0, \\ 0 & x < 0. \end{cases}$$

The corresponding loglikelihood function is

$$\begin{aligned} l(\lambda, k; x_0, x_1, \dots, x_n) &= \sum_{i=1}^n \ln \left[\frac{k}{\lambda} \left(\frac{x_i}{\lambda}\right)^{k-1} \exp\left(-\left(\frac{x_i}{\lambda}\right)^k\right) \right] \\ &= n(\ln k - k \ln \lambda) + (k-1) \sum_{i=1}^n \ln x_i - \sum_{i=1}^n \left(\frac{x_i}{\lambda}\right)^k. \end{aligned}$$

Regarding the Gamma distribution, its probability density function is

$$f(x; \alpha, \beta) = \begin{cases} \frac{x^{\alpha-1} \exp(-\frac{x}{\beta})}{\beta^\alpha \Gamma(\alpha)} & x > 0, \\ 0 & x \leq 0, \end{cases}$$

where Γ is the gamma function. Its loglikelihood function is

$$\begin{aligned}
l(\alpha, \beta; x_0, x_1, \dots, x_n) &= \ln \left(\prod_{i=1}^n \frac{x_i^{\alpha-1} \exp\left(-\frac{x_i}{\beta}\right)}{\beta^\alpha \Gamma(\alpha)} \right) \\
&= (a-1) \sum_{i=1}^n \ln x_i - \frac{1}{\beta} \sum_{i=1}^n x_i - n(\alpha \ln \beta + \ln \Gamma(\alpha)).
\end{aligned}$$

For the generalized Pareto distribution, it has the probability density function

$$f(x|k, \sigma, x_0) = \begin{cases} \frac{1}{\sigma} \left(1 + \frac{k(x-x_0)}{\sigma}\right)^{-1-1/k} & ; x > x_0 \text{ and } k > 0 \\ \frac{1}{\sigma} \exp\left(-\frac{x-x_0}{\sigma}\right) & ; x > x_0 \text{ and } k = 0 \\ \frac{1}{\sigma} \left(1 + \frac{k(x-x_0)}{\sigma}\right)^{-1-1/k} & ; x_0 < x < x_0 - \frac{\sigma}{k} \text{ and } k < 0. \end{cases}$$

For the simplicity, assume that $x_0 = 0$ and let $x^* = \max\{x_0, x_1, \dots, x_n\}$. The loglikelihood function is

$$l(\sigma, k; x_1, \dots, x_n) = \begin{cases} -n \ln \sigma + \left(\frac{1}{k} - 1\right) \sum_{i=1}^n \ln \left(1 - \frac{kx_i}{\sigma}\right), & k \neq 0 \\ -n \ln \sigma - \frac{1}{\sigma} \sum_{i=1}^n x_i, & k = 0, \end{cases}$$

where $\sigma > 0$ for $k \leq 0$ and $\sigma > k x^*$ for $k > 0$. The derivation of $l(\sigma, k; x_1, \dots, x_n)$ can be found in DuMouchel (1984) and Joe (1987). The optimization algorithm of $l(\sigma, k; x_1, \dots, x_n)$ can be found in D.Grimshaw (1993).

The following Matlab function carries out the Monte-Carlo simulation for giving critical values for EDF statistics. It is used to test the fit of duration distribution using exponential distributions and Weibull distributions. The code for Gamma distributions and generalized Pareto distribution is basically the same. The differences are the random number generating function and format of inputs. It consists of two subfunctions: stacal and critical that are presented later.

```
function [ exptable,wbltable ] = staMC(para,N,M,alpha )
% this function executes the mc
% para is the parameters, N is the length of data(without zeros),M is the
% size of simulation,alpha is the significance
format long
for i=1:M
exp=exprnd(para(1),para(2),1,N); % The random numbers for exp distribution
expx=sort(exp);
expz=cdf('exp',expx,para(1),para(2));
expresult(i,:)=stacal(expz);

wbl=wblrnd(para(3),para(4),1,N); % The random number matrix for wbl distribution
wblx=sort(wbl);
wblz=cdf('wbl',wblx,para(3),para(4));
wblresult(i,:)=stacal(wblz);
end

exptable=critical(expresult,alpha);
wbltable=critical(wblresult,alpha);
end
```

The Matlab function stacal is used to calculate the EDF statistics.

```
function [ result ] =stacal(z)
% since log, values in z must be positive
n=length(z);
for i=1:n
Dp(i)=(i/n-z(i));
Dm(i)=(z(i)-(i-1)/n);
v1(i)=(z(i)-(2*i-1)/(2*n))^2;
v2(i)=(2*i-1)*(log(z(i))+log(1-z(n+1-i)));
end
D1=max(Dp);

D2=max(Dm);

D=max(D1,D2);

W=sum(v1)+1/(12*n);

V=max(Dp)+max(Dm);

U=W-n*(sum(z)/n-1/2)^2;
A=-sum(v2)/n-n;

result=[D,V,W,U,A];
end
```

The Matlab function critical is used to give critical values.

```
function [ cvalue ] = critical(stat,alpha)
% stats is the statistics and alpha is the significance
[m,n]=size(stat);
k=length(alpha);
for i=1:n
x=sort(stat(:,i));
for j=1:k
index=round(m*(1-alpha(j)));
cvalue(i,j)=x(index);
end
end
end
```

References

- Admati, A. and Pfleiderer, P., 1988. A Theory of Intraday Patterns: Volume and Price Variability. *Review of Financial Studies*, 1(1), pp.3-40.
- Admati, A. and Pfleiderer, P., 1988. Selling and trading on information in financial markets. *American Economic Review*, 78 (1988), pp. 96-103
- Andersen, T. and Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4(2-3), pp.115-158.
- Andersen, T. and Bollerslev, T., 1998. Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies. *The Journal of Finance*, 53(1), pp.219-265.
- Andersen, T., 1996. Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility. *The Journal of Finance*, 51(1), pp.169-204.
- Andersen, T. G. (1994). Stochastic Autoregressive Volatility: A Framework for Volatility Modeling. *Mathematical Finance*, 4, 75-102.
- Andersen, T., 2001. The distribution of realized stock return volatility. *Journal of Financial Economics*, 61(1), pp.43-76.
- Andersen, T., Bollerslev, T., Diebold, F. and Labys, P., 2000. Exchange Rate Returns Standardized by Realized Volatility are (Nearly) Gaussian. *Multinational Finance Journal*, 4(3/4), pp.159-179.
- Andersen, T., T Bollerslev, FX Diebold, P Labys 1999. Realized volatility and correlation. LN Stern School of Finance Department Working Paper.
- Ané, T. and Geman, H., 2000. Order Flow, Transaction Clock, and Normality of Asset Returns. *The Journal of Finance*, 55(5), pp.2259-2284.

Baillie, R. and Bollerslev, T., 1991. Intra-Day and Inter-Market Volatility in Foreign Exchange Rates. *The Review of Economic Studies*, 58(3), p.565.

Barclay, M. and Warner, J., 1993. Stealth trading and volatility. *Journal of Financial Economics*, 34(3), pp.281-305.

Bauwens and Giot, 2000. The Logarithmic ACD Model: An Application to the Bid-Ask Quote Process of Three NYSE Stocks. *Annales d'Économie et de Statistique*, (60), p.117.

Bauwens, L. and Giot, P., 2003. Asymmetric ACD models: Introducing price information in ACD models. *Empirical Economics*, 28(4), pp.709-731.

Bauwens, L. and Veredas, D., 2004. The stochastic conditional duration model: a latent variable model for the analysis of financial durations. *Journal of Econometrics*, 119(2), pp.381-412.

Bauwens, L., Galli, F. and Giot, P., 2003. The Moments of Log-ACD Models. *SSRN Electronic Journal*,.

Bauwens, L., Giot, P., Grammig, J. and Veredas, D., 2004. A comparison of financial duration models via density forecasts. *International Journal of Forecasting*, 20(4), pp.589-609.

Bauwens, L., Preminger, A. and Rombouts, J., 2006. Regime Switching GARCH Models. *SSRN Electronic Journal*,.

Beltratti, A. and Morana, C., 1999. Computing value at risk with high frequency data. *Journal of Empirical Finance*, 6(5), pp.431-455.

Bertram, W., 2004. An empirical investigation of Australian Stock Exchange data. *Physica A: Statistical Mechanics and its Applications*, 341, pp.533-546.

Blume, L., Easley, D. and O'hara, M., 1994. Market Statistics and Technical Analysis: The Role of Volume. *The Journal of Finance*, 49(1), pp.153-181.

- Blume, L., Easley, D. and O'Hara, M., 1994. Market Statistics and Technical Analysis: The Role of Volume. *The Journal of Finance*, 49(1), pp.153-181.
- Bollerslev, T. and Domowitz, I., 1993. Trading Patterns and Prices in the Interbank Foreign Exchange Market. *The Journal of Finance*, 48(4), pp.1421-1443.
- Bollerslev, T. and Ghysels, E., 1996. Periodic Autoregressive Conditional Heteroscedasticity. *Journal of Business & Economic Statistics*, 14(2), p.139.
- Bollerslev, T. and Ole Mikkelsen, H., 1996. Modeling and pricing long memory in stock market volatility. *Journal of Econometrics*, 73(1), pp.151-184.
- Bollerslev, T. and Wright, J., 2000. High Frequency Data, Frequency Domain Inference and Volatility Forecasting. *SSRN Electronic Journal*,.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3), pp.307-327.
- Bollerslev, T., Chou, R. and Kroner, K., 1992. ARCH modeling in finance. *Journal of Econometrics*, 52(1-2), pp.5-59.
- Campbell, J., Grossman, S. and Wang, J., 1993. Trading Volume and Serial Correlation in Stock Returns. *The Quarterly Journal of Economics*, 108(4), pp.905-939.
- Carrasco, M. and Chen, X., 2002. Mixing and Moments Properties of Various GARCH and Stochastic Volatilities Models. *Econometric Theory*, 18(1), pp.17-39.
- Chanda, A., Engle, R. and Sokalska, M., 2005. High Frequency Multiplicative Component GARCH. *SSRN Electronic Journal*,.
- Chernov, M. and Ghysels, E., 2000. A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics*, 56(3), pp.407-458.

Christoffersen, P. and Pelletier, D., 2003. Backtesting Value-at-Risk: A Duration-Based Approach. SSRN Electronic Journal,.

Clark, P., 1973. A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices. *Econometrica*, 41(1), p.135.

Cont, R. and Bouchaud, J., 1998. Herd Behavior and Aggregate Fluctuations in Financial Markets. SSRN Electronic Journal,.

Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2), pp.223-236.

Coronel-Brizio, H. and Hernández-Montoya, A., 2005. On fitting the Pareto–Levy distribution to stock market index data: Selecting a suitable cutoff value. *Physica A: Statistical Mechanics and its Applications*, 354, pp.437-449.

Cox, D. and Isham, V., 2018. *Point Processes*. Boca Raton: Routledge.

Cumby, R., Figlewski, S. and Hasbrouck, J., 1993. Forecasting Volatilities and Correlations with EGARCH Models. *The Journal of Derivatives*, 1(2), pp.51-63.

Dacorogna, M., Müller, U., Nagler, R., Olsen, R. and Pictet, O., 1993. A geographical model for the daily and weekly seasonal volatility in the foreign exchange market. *Journal of International Money and Finance*, 12(4), pp.413-438.

Dacorogna, M., Müller, U., Pictet, O. and Olsen, R., 1997. Modelling Short-Term Volatility with GARCH and HARCH Models. SSRN Electronic Journal,.

Daley, D. and Vere-Jones, D., 2003. *An Introduction to The Theory of Point Processes*. New York: Springer.

Diamond, D. and V Verrecchia, R., 1991. Disclosure, Liquidity, and the Cost of Capital. *The Journal of Finance*, 46(4), pp.1325-1359.

Diamond, D. and Verrecchia, R., 1987. Constraints on short-selling and asset price adjustment to private information. *Journal of Financial Economics*, 18(2), pp.277-311.

- Diebold, F., Hahn, J. and Tay, A., 1999. Multivariate Density Forecast Evaluation and Calibration in Financial Risk Management: High-Frequency Returns on Foreign Exchange. *Review of Economics and Statistics*, 81(4), pp.661-673.
- Drost, F. and Nijman, T., 1993. Temporal Aggregation of Garch Processes. *Econometrica*, 61(4), p.909.
- Drost, F. and Werker, B., 1996. Closing the GARCH gap: Continuous time GARCH modeling. *Journal of Econometrics*, 74(1), pp.31-57.
- Drost, F. and Werker, B., 2004. Semiparametric Duration Models. *Journal of Business & Economic Statistics*, 22(1), pp.40-50.
- Dufour, A. and Engle, R. F. (2000): The ACD model: Predictability of the time between consecutive trades. Working Paper, ISMA Centre, University of Reading.
- DuMouchel, W., 1983. Estimating the Stable Index α in Order to Measure Tail Thickness: A Critique. *The Annals of Statistics*, 11(4), pp.1019-1031.
- Easley, D. and O'Hara, M., 1992. Time and the Process of Security Price Adjustment. *The Journal of Finance*, 47(2), pp.577-605.
- Easley, D. and O'Hara, M., 1995, Market microstructure, *Handbooks in Operations Research and Management Science*, Vol.9, pp. 357-384.
- Easley, D., Kiefer, N., O'hara, M. and Paperman, J., 1996. Liquidity, Information, and Infrequently Traded Stocks. *The Journal of Finance*, 51(4), pp.1405-1436.
- Engle, R. and Gallo, G., 2006. A multiple indicators model for volatility using intra-daily data. *Journal of Econometrics*, 131(1-2), pp.3-27.
- Engle, R. and Russell, J., 1997. Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model. *Journal of Empirical Finance*, 4(2-3), pp.187-212.
- Engle, R. and Russell, J., 1998. Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data. *Econometrica*, 66(5), p.1127.

Engle, R. and Sokalska, M., 2011. Forecasting intraday volatility in the US equity market. Multiplicative component GARCH. *Journal of Financial Econometrics*, 10(1), pp.54-83.

Engle, R., 1982. A general approach to lagrange multiplier model diagnostics. *Journal of Econometrics*, 20(1), pp.83-104.

Engle, R., 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50(4), p.987.

Engle, R., 2000. The Econometrics of Ultra-high-frequency Data. *Econometrica*, 68(1), pp.1-22.

Engle, R., 2002. Dynamic Conditional Correlation. *Journal of Business & Economic Statistics*, 20(3), pp.339-350.

Engle, R., Hendry, D. and Richard, J., 1983. Exogeneity. *Econometrica*, 51(2), p.277.

Fama, E., 1965. The Behavior of Stock-Market Prices. *The Journal of Business*, 38(1), p.34.

Farrell, L., 1996. Foreign Exchange Volume: Sound and fury, signifying nothing. *BMJ*, 313(7050), pp.174-174.

Feng, D., 2004. Stochastic Conditional Duration Models with "Leverage Effect" for Financial Transaction Data. *Journal of Financial Econometrics*, 2(3), pp.390-421.

Focardi, S. and Fabozzi, F., 2005. An autoregressive conditional duration model of credit - risk contagion. *The Journal of Risk Finance*, 6(3), pp.208-225.

Gallant, A., Hsu, C. and Tauchen, G., 2000. Using Daily Range Data to Calibrate Volatility Diffusions and Extract the Forward Integrated Variance. *SSRN Electronic Journal*,.

Gallant, A., Rossi, P.E., Tauchen G., 1993. Nonlinear dynamic structures. *Econometrica*, 61, pp. 871-907

- Gallant, A., Rossi, P. and Tauchen, G., 1992. Stock Prices and Volume. *Review of Financial Studies*, 5(2), pp.199-242.
- Gerhard, F. and Hautsch, N., 2001. Volatility Estimation on the Basis of Price Intensities. *SSRN Electronic Journal*,.
- Ghashghaie, S., Breymann, W., Peinke, J., Talkner, P. and Dodge, Y., 1996. Turbulent cascades in foreign exchange markets. *Nature*, 381(6585), pp.767-770.
- Ghose D., K.F. Kroner. 1994. Temporal aggregation of high frequency data Department of Economics, University of Arizona, Unpublished manuscript.
- Ghose, D. and Kroner, K., 1995. The relationship between GARCH and symmetric stable processes: Finding the source of fat tails in financial data. *Journal of Empirical Finance*, 2(3), pp.225-251.
- Ghysels, E., and J. Jasiak, 1998, GARCH for irregularly spaced financial data: The ACDGARCH model, *Studies in Nonlinear Dynamics and Econometrics*, 2, 133–149.
- Ghysels, E., C. Gouriéroux, and J. Jasiak, 1998, Stochastic volatility duration models, CIRANO.
- Ghysels, E., Lee, H. and Siklos, P., 1993. On the (mis)specification of seasonality and its consequences: An empirical investigation with US data. *Empirical Economics*, 18(4), pp.747-760.
- Giot, P., 2003. Implied Volatility Indices as Leading Indicators of Stock Index Returns?. *SSRN Electronic Journal*,.
- Giot, P., 2005. Market risk models for intraday data. *The European Journal of Finance*, 11(4), pp.309-324.
- Gnaciński, P. and Makowiec, D., 2004. Another type of log-periodic oscillations on Polish stock market. *Physica A: Statistical Mechanics and its Applications*, 344(1-2), pp.322-325.

Goodhart, C. and O'Hara, M., 1997. High frequency data in financial markets: Issues and applications. *Journal of Empirical Finance*, 4(2-3), pp.73-114.

Gopikrishnan, P., Meyer, M., Amaral, L. and Stanley, H., 1998. Inverse cubic law for the distribution of stock price variations. *The European Physical Journal B*, 3(2), pp.139-140.

Gramming J., Maurer K.O. (1999) Non-monotonic hazard functions and the autoregressive conditional duration model, Discussion Paper 50, SFB 373, Humboldt University Berlin.

Gramming, J., and Wellner, M. (2002). 'Modeling the interdependence of volatility and intertransaction duration processes', *Journal of Econometrics*, 106: 369 400.

Gu, G., Chen, W. and Zhou, W., 2007. Quantifying bid-ask spreads in the Chinese stock market using limit-order book data. *The European Physical Journal B*, 57(1), pp.81-87.

Guillaume, D., Dacorogna, M., Davé, R., Müller, U., Olsen, R. and Pictet, O., 1997. From the bird's eye to the microscope: A survey of new stylized facts of the intra-daily foreign exchange markets. *Finance and Stochastics*, 1(2), pp.95-129.

H. McNish, T. and Wood, R., 1990. A transactions data analysis of the variability of common stock returns during 1980–1984. *Journal of Banking & Finance*, 14(1), pp.99-112.

Harris, L. and Gurel, E., 1986. Price and Volume Effects Associated with Changes in the S&P 500 List: New Evidence for the Existence of Price Pressures. *The Journal of Finance*, 41(4), p.815.

Harris, L., 1987. Transaction Data Tests of the Mixture of Distributions Hypothesis. *The Journal of Financial and Quantitative Analysis*, 22(2), p.127.

Hasbrouck, J. and Sofianos, G., 1993. The Trades of Market Makers: An Empirical Analysis of NYSE Specialists. *The Journal of Finance*, 48(5), pp.1565-1593.

Hasbrouck, J., 1991. The Summary Informativeness of Stock Trades: An Econometric Analysis. *Review of Financial Studies*, 4(3), pp.571-595.

- Hautsch, N., 2004. *Modelling Irregularly Spaced Financial Data*. Berlin: Springer.
- Heston, S., 1993. Invisible Parameters in Option Prices. *The Journal of Finance*, 48(3), pp.933-947.
- Hörmann, S. (2009). Berry-Esseen bounds for econometric time series. *ALEA* 6 377–397.
- Hujer, R. and Vuletic, S., 2005. The Regime Switching ACD Framework: The Use of the Comprehensive Family of Distributions. *SSRN Electronic Journal*,.
- Jansen, D. and de Vries, C., 1991. On the Frequency of Large Stock Returns: Putting Booms and Busts into Perspective. *The Review of Economics and Statistics*, 73(1), p.18.
- Jasiak, J., 1999. Persistence in Intertrade Durations. *SSRN Electronic Journal*,.
- Jirak, M., 2016. Berry–Esseen theorems under weak dependence. *The Annals of Probability*, 44(3), pp.2024-2063.
- Jones, C., Kaul, G. and Lipson, M., 1994. Transactions, Volume, and Volatility. *Review of Financial Studies*, 7(4), pp.631-651.
- Karpoff, J., 1987. The Relation Between Price Changes and Trading Volume: A Survey. *The Journal of Financial and Quantitative Analysis*, 22(1), p.109.
- Lamoureux, C. and Lastrapes, W., 1990. Persistence in Variance, Structural Change, and the GARCH Model. *Journal of Business & Economic Statistics*, 8(2), pp.225-234.
- Lancaster, T., n.d. *The Econometric Analysis of Transition Data*. Florida: Chapman&Hall/CRC
- Lee, S. and Hansen, B., 1994. Asymptotic Theory for the Garch (1,1) Quasi-Maximum Likelihood Estimator. *Econometric Theory*, 10(1), pp.29-52.
- Lehmann, B. and Modest, D., 1994. Trading and Liquidity on the Tokyo Stock Exchange: A Bird's Eye View. *The Journal of Finance*, 49(3), pp.951-984.

Lin, S. and Tamvakis, M., 2004. Effects of NYMEX trading on IPE Brent Crude futures markets: a duration analysis. *Energy Policy*, 32(1), pp.77-82.

Ling, S. and Li, W., 1997. On Fractionally Integrated Autoregressive Moving-Average Time Series Models with Conditional Heteroscedasticity. *Journal of the American Statistical Association*, 92(439), pp.1184-1194.

Ling, S. and Li, W., 1998. Limiting distributions of maximum likelihood estimators for unstable autoregressive moving-average time series with general autoregressive heteroscedastic errors. *The Annals of Statistics*, 26(1), pp.84-125.

Ling, S. and McAleer, M., 2003. Asymptotic theory for a vector ARMA-GARCH model. *Econometric Theory*, 19(02).

Lockwood, L. and Linn, S., 1990. An Examination of Stock Market Return Volatility During Overnight and Intraday Periods, 1964-1989. *The Journal of Finance*, 45(2), pp.591-601.

Loretan, M. and Phillips, P., 1994. Testing the covariance stationarity of heavy-tailed time series: An overview of the theory with applications to several financial datasets. *Journal of Empirical Finance*, 1(2), pp.211-248.

Lumsdaine, R., 1996. Consistency and Asymptotic Normality of the Quasi-Maximum Likelihood Estimator in IGARCH(1,1) and Covariance Stationary GARCH(1,1) Models. *Econometrica*, 64(3), p.575.

Lunde, A., 1999. A generalized Gamma autoregressive conditional duration model. Working Paper, Department of Economics, Politics and Public Administration, Aalborg University, Denmark.

Mandelbrot, B., 1963. The Stable Paretian Income Distribution when the Apparent Exponent is Near Two. *International Economic Review*, 4(1), p.111.

Manganelli, S., 2005. Duration, volume and volatility impact of trades. *Journal of Financial Markets*, 8(4), pp.377-399.

Mantegna, R. and Stanley, H., 1995. Scaling behaviour in the dynamics of an economic index. *Nature*, 376(6535), pp.46-49.

McCulloch, J., 1997. Measuring Tail Thickness to Estimate the Stable Index α : A Critique. *Journal of Business & Economic Statistics*, 15(1), p.74.

Meitz, M. and Teräsvirta, T., 2006. Evaluating Models of Autoregressive Conditional Duration. *Journal of Business & Economic Statistics*, 24(1), pp.104-124.

Mittink, S. and Rachev, S., 1993. Reply to comments on modeling asset returns with alternative stable distributions and some extensions*. *Econometric Reviews*, 12(3), pp.347-389.

Müller U., Darorogna, M., Davé R., Pictet O, Olsen R., Ward J. 1993. Fractals and intrinsic time: A challenge to econometricians, Olsen and Associates, Research Institute for Applied Economics, Working Paper.

Müller, U., Dacorogna, Davé R., Olsen R., O., Pictet, O. 1995. Volatilities of different time resolutions: Analyzing the dynamics of market components, Olsen and Associates, Research Institute for Applied Economics, Working Paper.

Müller, U., Dacorogna, M., Davé, R., Olsen, R., Pictet, O. and von Weizsäcker, J., 1997. Volatilities of different time resolutions — Analyzing the dynamics of market components. *Journal of Empirical Finance*, 4(2-3), pp.213-239.

Müller, U., Dacorogna, M., Olsen, R., Pictet, O., Schwarz, M. and Morgenegg, C., 1990. Statistical study of foreign exchange rates, empirical evidence of a price change scaling law, and intraday analysis. *Journal of Banking & Finance*, 14(6), pp.1189-1208.

Nelson, D. and Cao, C., 1992. Inequality Constraints in the Univariate GARCH Model. *Journal of Business & Economic Statistics*, 10(2), p.229.

Nelson, D., 1991. Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica*, 59(2), p.347.

Newey, W., 1994. The Asymptotic Variance of Semiparametric Estimators. *Econometrica*, 62(6), p.1349.

Nolan, J., 2018. Stable Distributions and Green's Functions for Fractional Diffusions. *SSRN Electronic Journal*,.

Pacurar, M., 2006. Autoregressive Conditional Duration (ACD) Models in Finance: A Survey of the Theoretical and Empirical Literature. *SSRN Electronic Journal*,.

Prigent, J., 2001. Option Pricing with a General Marked Point Process. *Mathematics of Operations Research*, 26(1), pp.50-66.

Richardson, M. and Smith, T., 1994. A Direct Test of the Mixture of Distributions Hypothesis: Measuring the Daily Flow of Information. *The Journal of Financial and Quantitative Analysis*, 29(1), p.101.

Ross, S., 1989. Information and Volatility: The No-Arbitrage Martingale Approach to Timing and Resolution Irrelevancy. *The Journal of Finance*, 44(1), pp.1-17.

Russell, J. and Engle, R., 1998. Econometric Analysis of Discrete-Valued Irregularly-Spaced Financial Transactions Data Using a New Autoregressive Conditional Multinomial Model. *SSRN Electronic Journal*,.

Russell, J. and Engle, R., 2005. A Discrete-State Continuous-Time Model of Financial Transactions Prices and Times. *Journal of Business & Economic Statistics*, 23(2), pp.166-180.

Scalas, E., Gorenflo, R. and Mainardi, F., 2000. Fractional calculus and continuous-time finance. *Physica A: Statistical Mechanics and its Applications*, 284(1-4), pp.376-384.

Spierdijk, L., 2004. An empirical analysis of the role of the trading intensity in information dissemination on the NYSE. *Journal of Empirical Finance*, 11(2), pp.163-184.

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

Taylor, S. and Xu, X., 1997. The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance*, 4(4), pp.317-340.

Tsay, R., 2003. *Analysis of Financial Time Series*. Wiley.

Veredas, D., Rodriguez-Poo, J., & Espasa, A. (2002). On the (intradaily) seasonality and dynamics of a financial point process: A semiparametric approach, CORE Discussion Paper, 23, Université Catholique de Louvain.

West, K. and Cho, D., 1995. The predictive ability of several models of exchange rate volatility. *Journal of Econometrics*, 69(2), pp.367-391.

Wood, R., Mcinish, T. and ORD, J., 1985. An Investigation of Transactions Data for NYSE Stocks. *The Journal of Finance*, 40(3), pp.723-739.

Whittle, P., 1953. The Analysis of Multiple Stationary Time Series. *Journal of the Royal Statistical Society: Series B (Methodological)*, 15(1), pp.125-139.

Young, M. and Graff, R., 1995. Real estate is not normal: A fresh look at real estate return distributions. *The Journal of Real Estate Finance and Economics*, 10(3), pp.225-259.

Zhang, L., 2010. Implied and realized volatility: empirical model selection. *Annals of Finance*, 8(2-3), pp.259-275.

Zhang, L., Mykland, P. and Aït-Sahalia, Y., 2005. A Tale of Two Time Scales. *Journal of the American Statistical Association*, 100(472), pp.1394-1411.