

Approaches to Modeling Market Liquidity

by

Annie Shoup

B.S., Massachusetts Institute of Technology (2016)

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2019

© Massachusetts Institute of Technology 2019. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 24, 2019

Certified by.....
Dr. Peter Kempthorne
Department of Mathematics
Thesis Supervisor

Accepted by.....
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Approaches to Modeling Market Liquidity

by

Annie Shoup

Submitted to the Department of Electrical Engineering and Computer Science
on May 24, 2019, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Intraday data was collected for the U.S. Financial Sector ETF and two of its component stocks: Bank of America and Citigroup. The analysis period includes 32 trading days, ranging from February 4, 2019 to March 20, 2019. From this trade and quote data, we construct 2,496 five-minute aggregate time bars for each security and calculate a series of spread, volume, depth, trade count, and price change liquidity measures. We examine the summary statistics of these liquidity measures before applying principal components analysis to them through which we identify key liquidity dimensions in each security and common liquidity dimensions across them. Vector autoregressive models are applied to these principal component scores in order to gain further insight into their time series structure and the ways in which the measures interact over a 32-day period. Finally, the same methodology of principal components and time series analyses are applied to daily-normalized liquidity measures in order to better understand the intraday, rather than multi-day, dynamics of liquidity.

Thesis Supervisor: Dr. Peter Kempthorne

Acknowledgments

Dr. Kempthorne, thank you for your continued support and guidance throughout this process. This thesis was a truly invaluable experience from which I learned an incredible amount.

Samira and Andrew, thank you for always providing me a comfortable couch to work from.

Emma, thank you for providing me with multiple comfortable couches to work from; the leather sectional in your living room is entirely responsible for Chapter 2 of this thesis.

Allison, Alex, Isabella, and Lexi, while you didn't provide me any couches to work from (hence being mentioned fourth), thank you for the constant encouragement and optimism that you did provide me with throughout this process.

Finally, thank you to Ben (and Pilgrim). I would not have been able to finish this without you.

Contents

1	Introduction	25
2	Liquidity Measures for Equity Markets	29
2.1	Bid-ask Spreads	29
2.2	Volume and Depth Measures	37
2.3	Trade Count Measures	39
2.4	Price Change Measures	40
3	Transactions Data for the Financial Sector ETF and Two Component Stocks	43
3.1	Financial Sector Exchange-Traded Fund (XLF)	45
3.2	Exchange-Traded Fund (ETF) Market Mechanics	47
3.3	ETF Asset Valuation	48
4	Summary Statistics of Liquidity Measures	51
4.1	Liquidity Measures for Bank of America	53
4.2	Liquidity Measures for Citigroup	63
4.3	Liquidity Measures for the Financial Sector ETF	76
5	Principal Components Analysis of Liquidity Measures	91
5.1	Principal Components Analysis Theory	92
5.2	Varimax Rotation of Principal Component Loadings	94
5.3	Principal Components Analysis of Bank of America's Liquidity Measures	95
5.4	Principal Components Analysis of Citigroup's Liquidity Measures . . .	100

5.5	Principal Components Analysis of the Financial Sector ETF's Liquidity Measures	105
6	Time Series Analysis of Principal Component Scores of Liquidity Measures	111
6.1	Autocorrelation and Cross-correlation in Liquidity Measures	112
6.2	Vector Autoregressive Model	113
6.3	Bank of America Time Series Analysis	114
6.3.1	Time Series Structure	115
6.3.2	Vector Autoregressive Model	123
6.4	Citigroup Time Series Analysis	128
6.4.1	Time Series Structure	130
6.4.2	Vector Autoregressive Model	133
6.5	Financial Sector ETF Time Series Analysis	136
6.5.1	Time Series Structure	137
6.5.2	Vector Autoregressive Model	138
7	Principal Components Analysis of Daily-normalized Liquidity Measures	145
7.1	Principal Components Analysis of Bank of America's Daily-normalized Liquidity Measures	146
7.2	Principal Component Analysis of Citigroup's Daily-normalized Liquidity Measures	151
7.3	Principal Components Analysis of the Financial Sector ETF's Daily-normalized Liquidity Measures	157
8	Time Series Analysis of Principal Component Scores of Daily-normalized Liquidity Measures	163
8.1	Bank of America Time Series Analysis	163
8.1.1	Time Series Structure	164
8.1.2	Vector Autoregressive Model	166

8.2	Citigroup Time Series Analysis	171
8.2.1	Time Series Structure	172
8.2.2	Vector Autoregressive Model Regressive Model	172
8.3	Financial Sector ETF Time Series Analysis	177
8.3.1	Time Series Structure	177
8.3.2	Vector Autoregressive Model	179
9	Conclusion	185

List of Figures

4-1	The average, standard deviation, and coefficient of variation of Bank of America’s spread measures, aggregated across 2,483 five-minute bars.	56
4-2	The average, standard deviation, and coefficient of variation of Bank of America’s volume and depth measures, aggregated across 2,483 five-minute bars.	58
4-3	The average, standard deviation, and coefficient of variation of Bank of America’s trade count measures, aggregated across 2,483 five-minute bars.	60
4-4	The average, standard deviation, and coefficient of variation of Bank of America’s price-change measures, aggregated across 2,483 five-minute bars.	61
4-5	The Z-scores of Bank of America’s liquidity measures averaged across 15-minute intervals.	62
4-6	The average, standard deviation, and coefficient of variation of Citigroup’s spread measures, aggregated across 2,485 five-minute bars. . .	64
4-7	The average, standard deviation, and coefficient of variation of Citigroup’s volume and depth measures, aggregated across 2,485 five-minute bars.	67
4-8	The average, standard deviation, and coefficient of variation of Citigroup’s trade count measures, aggregated across 2,485 five-minute bars.	69
4-9	The average, standard deviation, and coefficient of variation of Citigroup’s price-change measures, aggregated across 2,485 five-minute bars.	71

4-10	The Z-scores of Citigroup’s liquidity measures averaged across 15-minute intervals.	72
4-11	The Z-scores of Citigroup’s spread measures averaged across 15-minute intervals.	73
4-12	The Z-scores of Citigroup’s volume and depth measures averaged across 15-minute intervals.	74
4-13	The Z-scores of Citigroup’s trade count measures averaged across 15-minute intervals.	75
4-14	The Z-scores of Citigroup’s price change composite measures averaged across 15-minute intervals.	76
4-15	The average, standard deviation, and coefficient of variation of the Financial Sector ETF’s spread measures, aggregated across 2,496 five-minute bars.	78
4-16	The average, standard deviation, and coefficient of variation of the Financial Sector ETF’s volume and depth measures, aggregated across 2,496 five-minute bars.	80
4-17	The average, standard deviation, and coefficient of variation of the Financial Sector ETF’s trade count measures, aggregated across 2,496 five-minute bars.	82
4-18	The average, standard deviation, and coefficient of variation of the Financial Sector ETF’s price-change measures, aggregated across 2,496 five-minute bars.	84
4-19	The Z-scores of the Financial Sector ETF’s liquidity measures averaged across 15-minute intervals.	85
4-20	The Z-scores of the Financial Sector ETF’s spread measures averaged across 15-minute intervals.	86
4-21	The Z-scores of the Financial Sector ETF’s volume and depth measures averaged across 15-minute intervals.	87
4-22	The Z-scores of the Financial Sector ETF’s trade count measures averaged across 15-minute intervals.	88

4-23	The Z-scores of the Financial Sector ETF's price change composite measures averaged across 15-minute intervals.	89
5-1	The scree plot of Bank of America's principal components resulting from PCA on 13 liquidity measures.	96
5-2	The principal component loadings of Bank of America resulting from PCA on 13 liquidity measures.	97
5-3	The scree plot of Citigroup's principal components resulting from PCA on 13 liquidity measures.	102
5-4	The principal component loadings of Citigroup resulting from PCA on 13 liquidity measures.	102
5-5	The scree plot of the Financial Sector ETF's principal components resulting from PCA on 13 liquidity measures.	107
5-6	The principal component loadings of the Financial Sector ETF resulting from PCA on 13 liquidity measures.	108
6-1	The cumulative sum of Bank of America's principal components scores over time.	115
6-2	The autocorrelation and cross-correlation of Bank of America's principal component scores.	117
6-3	The autocorrelation and partial autocorrelation of Bank of America's first principal component's scores.	118
6-4	The autocorrelation and partial autocorrelation of Bank of America's second principal component's scores.	119
6-5	The autocorrelation and partial autocorrelation of Bank of America's third principal component's scores.	120
6-6	The autocorrelation and partial autocorrelation of Bank of America's fourth principal component's scores.	121
6-7	The autocorrelation and partial autocorrelation of Bank of America's fifth principal component's scores.	122

6-8	The partial autocorrelation of Bank of America’s principal component scores.	123
6-9	The cumulative sum of Citigroup’s principal components scores over time.	129
6-10	The autocorrelation and cross-correlation of Citigroup’s principal components scores over time.	131
6-11	The partial autocorrelation and partial cross-correlation of Citigroup’s principal components scores over time.	132
6-12	The cumulative sum of the Financial Sector ETF’s significant principal components’ scores over time.	137
6-13	The autocorrelation and cross-correlation of the Financial Sector ETF’s principal components scores over time.	139
6-14	The partial autocorrelation and partial cross-correlation of the Financial Sector ETF’s principal components scores over time.	140
7-1	The scree plot of Bank of America’s principal components resulting from two principal component analyses; one on daily-normalized measures and the other on measures normalized over a 32-day time period.	147
7-2	Bank of America’s principal component loadings from principal components analysis on 13 daily-normalized liquidity measures.	148
7-3	The scree plot of Citigroup’s principal components resulting from two principal component analyses on 13 liquidity measures normalized daily and normalized over the entire 32-day time period.	152
7-4	The principal component loadings of Citigroup resulting from PCA on 13 liquidity measures.	153
7-5	The scree plots of the Financial Sector ETF’s principal components resulting from two principal component analyses, one on 2,946 observations of 13 liquidity measures each, where one analysis normalizes the measures over the entire time period (32 days) and the other normalizes them daily.	158

7-6	The principal component loadings of the Financial Sector ETF resulting from PCA on 2,496 daily-normalized observations of 13 liquidity measures.	159
8-1	The cumulative sum of BAC's principal components scores over time.	165
8-2	The autocorrelation and cross-correlation of Bank of America's principal components scores of Daily-normalized Liquidity measures.	166
8-3	The partial autocorrelation of Bank of America's principal components scores of daily-normalized liquidity measures.	167
8-4	The cumulative sum of Citigroup's principal components scores over time.	171
8-5	The first principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	171
8-6	The autocorrelation and cross-correlation of Citigroup's principal components scores of daily-normalized liquidity measures.	173
8-7	The partial autocorrelation of Citigroup's principal components scores of daily-normalized liquidity measures.	174
8-8	The cumulative sum of the Financial Sector ETF's principal components scores over time.	177
8-9	The autocorrelation and cross-correlation of the Financial Sector ETF's principal components scores of daily-normalized liquidity measures.	178
8-10	The partial autocorrelation of the Financial Sector ETF's principal components scores of daily-normalized liquidity measures.	180

List of Tables

2.1	A brief summary of the spread measures discussed where p is the asset price, a is the ask side, b is the bid side, m is the mid, t indicates an executed trade, and n is an amount of time over which price impact is measured.	36
3.1	The largest constituents that make up the Financial Select Sector SPDR Fund as of March 22, 2019. Only Bank of America and Citigroup are included in this analysis.	46
3.2	The top 10 holdings by market value of XLF as of March 22, 2019.	46
4.1	Summary statistics of Bank of America’s spread measures; these encompass 2,483 five-minute bar observations for each measure.	54
4.2	Summary statistics of Bank of America’s volume and depth measures; these encompass 2,483 five-minute bar observations for each measure.	57
4.3	Summary statistics of Bank of America’s trade count measures; these encompass 2,483 five-minute bar observations for each measure.	59
4.4	Summary statistics of Bank of America’s liquidity measures involving price changes; these encompass 2,483 five-minute bar observations for each measure.	60
4.5	Summary statistics of Citigroup’s spread measures; these encompass 2,485 five-minute bar observations for each measure.	65
4.6	Summary statistics of Citigroup’s volume and depth measures; these encompass 2,485 five-minute bar observations for each measure.	66

4.7	Summary statistics of Citigroup’s trade count measures; these encompass 2,485 five-minute bar observations for each measure.	68
4.8	Summary statistics of Citigroup’s liquidity measures involving price changes; these encompass 2,485 five-minute bar observations for each measure.	69
4.9	Summary statistics of the Financial Sector ETF’s spread measures; these encompass 2,496 five-minute bar observations for each measure.	77
4.10	Summary statistics of the Financial Sector ETF’s volume and depth measures; these encompass 2,496 five-minute bar observations for each measure.	79
4.11	Summary statistics of the Financial Sector ETF’s trade count measures; these encompass 2,496 five-minute bar observations for each measure.	81
4.12	Summary statistics of the Financial Sector ETF’s liquidity measures involving price changes; these encompass 2,496 five-minute bar observations for each measure.	83
5.1	The principal components of Bank of America using PCA on 13 liquidity measures and the Kaiser criterion to determine significance. . .	96
5.2	The principal component loadings of each of Bank of America’s 13 liquidity measures across components.	97
5.3	A varimax rotation of Bank of America’s first five principal component loadings resulting from principal components analysis on liquidity measures.	99
5.4	Percentage of cumulative variance of each of Bank of America’s original liquidity measures as explained by the five significant PCA components.	101
5.5	The principal components of Citigroup using PCA on 13 liquidity measures and the Kaiser criterion to determine significance.	101
5.6	The principal component loadings of Citigroup using PCA on 13 liquidity measures and the Kaiser criterion to determine significance. . .	104

5.7	A varimax rotation of Citigroup’s first four principal component loadings resulting from principal components analysis on liquidity measures.	104
5.8	The variance of Citigroup’s original variables explained by the significant principal components.	105
5.9	The principal components of the Financial Sector ETF using PCA on 13 liquidity measures and the Kaiser criterion to determine significance.	106
5.10	The principal component loadings of the Financial Sector ETF using PCA on 13 liquidity measures and the Kaiser criterion to determine significance.	109
5.11	A varimax rotation of the Financial Sector ETF’s first five principal component loadings resulting from principal components analysis on liquidity measures.	109
5.12	The variance of the Financial Sector ETF’s original variables explained by its five significant principal components.	110
6.1	The first principal component’s results from a second-order VAR model applied to Bank of America’s principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	124
6.2	The second principal component’s results from a second-order VAR model applied to Bank of America’s principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	125
6.3	The third principal component’s results from a second-order VAR model applied to Bank of America’s principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	125

6.4	The fourth principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	127
6.5	The fifth principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	128
6.6	The first principal component's results from a second-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	133
6.7	The second principal component's results from a second-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	134
6.8	The third principal component's results from a second-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	135
6.9	The fourth principal component's results from a second-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	136
6.10	The first principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	141

6.11	The second principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	142
6.12	The third principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	142
6.13	The fourth principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	143
6.14	The fifth principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	143
7.1	The first five principal components of a principal components analysis on Bank of America's 13 daily-normalized liquidity measures.	146
7.2	The principal component loadings of each of Bank of America's 13 daily-normalized liquidity measures across components.	149
7.3	A varimax rotation of Bank of America's first five principal component loadings resulting from principal components analysis on daily-normalized liquidity measures.	150
7.4	Percentage of cumulative variance of each of BAC's original liquidity measures as explained by the five significant PCA components.	151
7.5	The principal components of Citigroup using PCA on 13 daily-normalized liquidity measures and the Kaiser criterion to determine significance.	151
7.6	The principal component loadings of each of Citigroup's 13 liquidity measures across components.	154

7.7	A varimax rotation of Citigroup’s first four principal component loadings resulting from principal components analysis on daily-normalized liquidity measures.	155
7.8	Percentage of cumulative variance of each of Citigroup’s daily liquidity measures as explained by the four significant PCA components. . . .	156
7.9	The principal components of the Financial Sector ETF resulting from principal components analysis on 2,496 observations of 13 daily-normalized liquidity measures.	157
7.10	The principal component loadings of each of the Financial Sector ETF’s 13 liquidity measures across components.	160
7.11	A varimax rotation of the Financial Sector ETF’s first five principal component loadings resulting from principal components analysis on daily-normalized liquidity measures.	161
7.12	Percentage of cumulative variance of each of the Financial Sector ETF’s daily-normalized liquidity measures as explained by principal components analysis.	162
8.1	The first principal component’s results from a second-order VAR model applied to Bank of America’s principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	168
8.2	The second principal component’s results from a second-order VAR model applied to Bank of America’s principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	168
8.3	The third principal component’s results from a second-order VAR model applied to Bank of America’s principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	169

8.4	The fourth principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	169
8.5	The fifth principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	170
8.6	The first principal component's results from a third-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	175
8.7	The second principal component's results from a third-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	175
8.8	The third principal component's results from a third-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	176
8.9	The fourth principal component's results from a third-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	176
8.10	The first principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	181

8.11	The second principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	181
8.12	The third principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	182
8.13	The fourth principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	182
8.14	The fifth principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.	183

Chapter 1

Introduction

The existence of an asset does not create a market; rather the liquidity of the asset does. In the most fundamental sense, liquidity can be thought of as the presence of a sufficient number of both buyers and sellers. Operationally, liquidity is the ability to buy or sell an asset without causing a drastic change in the asset's price. When there is an imbalance of buyers and sellers (or in the most severe cases an absence), the resulting illiquidity can cause increased transaction costs, price slippage, and difficulty or even inability to close out a position.

While price slippage and an inability to buy or sell are both natural results from a lack of supply, market participants more directly induce illiquidity's effects on transaction costs. Illiquidity affects transaction costs because it introduces more risk and uncertainty into the transaction for the market maker. This liquidity risk is the potential future inability to buy or sell the asset, especially during times of market stress. As a result, market makers charge a higher cost for trading less liquid assets; this is effected through a larger bid-ask spread. To compensate for the greater transactions costs, investors require higher returns on assets with lower market liquidity. Accurately modeling liquidity can reduce transaction costs and enable one to more effectively ascertain the true value of an asset.

The tension between illiquidity and efficient asset pricing is particularly prevalent in the context of today's exchange-traded fund (ETF) market. Although ETFs started

trading in the U.S. in 1993, their volume was modest until recently. The financial crisis fundamentally altered the behavior of market participants. Investors sought ways to diversify their risk while banks looked to offload physical assets from their balance sheets to reduce new capital charges and inventory and risk limits arising from increased regulation. ETFs offered a cheap, passively-managed way for investors and banks to accomplish their goals.

As the volume of ETFs traded in the market increased rapidly, their usage has evolved. Today, ETFs usually account for a quarter of the daily volume in the U.S. stock market. On some days, that volume can increase to nearly 40%. Because ETFs provide a quick and cost-efficient way to execute macro views, these sharp increases in volume are usually preceded by surprise events and policy decisions.

However, as ETF ownership grows, an increasing proportion of the outstanding shares for the underlying security is held in trust by the fund sponsor. Although these shares are available for trade as part of a basket transaction at the ETF-level, they are no longer available to traders who wish to transact on firm-specific information. Because of this, as ETFs become larger holders of a firm's shares, transaction costs for the underlying securities can increase. This increase in trading costs is associated with a decrease in available liquidity for the component securities owned by ETFs.

The extent to which market liquidity affects asset returns and pricing efficiency varies with respect to an asset's liquidity. In a very liquid market, buying and selling reasonable quantities of the asset will not greatly affect its price. However, in a very illiquid market, buying and selling reasonable quantities of an asset will cause a significant adverse change in the price of the asset. What constitutes a reasonable amount of an asset is all relative to the normal activity in that asset's market and is different for different types of assets.

Because of the decrease in liquidity and consequent increase in transaction costs, larger ETF ownership is accompanied by a decline in the pricing efficiency of the underlying component securities. Modeling the liquidity of an ETF using its underlying components has the potential to reduce this pricing efficiency friction by enabling an investor to arbitrage the differences between the ETF's liquidity and that of its

constituents.

Beyond arbitrage opportunities that may arise from liquidity modeling, analyzing liquidity patterns has potential to help regulators detect insider trading. When an investor possesses material, non-public information and wishes to trade on it, the investor becomes an informed trader. Sophisticated insiders will strategically time their trades to avoid detection, waiting for days when liquidity is high and their trades are harder to identify by market participants and regulators. However, when information is short-lived, and informed investors do not have the luxury to wait for days with high liquidity, nearly all of the proxies of informed trading are correlated with illegal insider trading. These proxies for informed trading are liquidity metrics that identify unusual increases in volumes and associated price changes. Using these liquidity measures, regulators may have the potential to better identify insider trading on time sensitive information.

While the act of insider trading is presumably rare, the ability to identify informed trading has uses beyond market regulation. Throughout the course of a trading day, market makers encounter a mixture of informed and uninformed market participants; liquidity measures offer market makers the possibility of being able to more easily separate the informed trades from the noise. Market makers, who are consequently liquidity providers, are susceptible to asymmetric information when trading against counterparties with better information. During normal conditions, a market maker may be able to reasonably make profit with both an equal volume of buys and sells. But when the order flow becomes too skewed, it can cause losses for these liquidity providers. In a scenario where investors only wish to transact on one side of the market, the order flow is considered toxic. For this reason, liquidity measures can be a predictor of extreme price movements, which has value for both liquidity providers and investors.

Although the usefulness of liquidity measures is clearly wide-ranging, this thesis will focus on identifying and modeling different dimensions of liquidity for three securities: Bank of America (BAC), Citigroup (C), and the Financial Select Sector SPDR Fund (XLF), which is the financial sector ETF of the Standard and Poors.

The analysis ranges from February 4, 2019 to March 20, 2019 and includes 32 days of trade and quote data during U.S. market hours.

Chapter 2 outlines previous liquidity measures studied in the literature. Chapter 3 describes the basic properties and structure of the studied data and the market mechanics of ETFs. Chapter 4 examines the summary statistics for liquidity measures of each security. Chapter 5 applies principal components analysis to the securities' liquidity measures and attempts to identify key liquidity properties from the newly defined principal components. Chapter 6 investigates the time series structure of the principal components and analyzes these patterns using vector autoregressive models. Finally, while the analysis in Chapters 5 and 6 focuses on capturing the interactions of liquidity measures over a 32-day period, Chapters 7 and 8 perform the same principal components and time series analysis on daily-normalized liquidity measures; this is done in order to examine the intraday dynamics of the liquidity measures.

Chapter 2

Liquidity Measures for Equity Markets

2.1 Bid-ask Spreads

Perhaps the most intuitive measure of liquidity is one that naturally arises in any quote-driven market: the bid-ask spread. The bid-ask spread is the difference between the bid, the price at which a market maker is willing to buy, and the ask, the price at which a market maker is willing to sell. The mid price is the average of the best bid (the highest of the market makers' bids) and the best ask (the lowest of the market makers' asks); the mid price is meant to reflect the fair price of the asset.

The bid-ask acts as a liquidity premium or concession that is paid to the market maker for providing immediate liquidity. Economically, one expects the ask price to be greater than or equal to the bid price and consequently the bid-ask spread to be positive. In certain situations, the bid and ask are equal (a locked market) or the bid is greater than the ask (a crossed market). Locked and crossed markets can arise from differences or lags in reporting and timing conventions of different exchanges or other structural inconsistencies. Regardless of the cause, locked and crossed markets are short-lived to the point of being almost instantaneous as market participants quickly take advantage of the arbitrage opportunity.

The nature of the bid-ask spread changes depending on the capacity in which the market maker is acting in. A market maker can act as either a principal or agent in a transaction. When acting as an agent, the market maker brokers a transaction between two separate counterparties and is the intermediary in their transaction. Because the market maker is only an arranger and does not take any inventory risk when executing the transaction, the bid-ask can be thought of as a broker's fee in these cases. Thus when this occurs, the spread is perhaps not so representative of the asset's liquidity risk.

However, when a market maker acts as a principal, the market maker buys or sells from his own inventory to the counterparty. If the price of the asset were to decline after buying and before being able to sell, the market maker would incur a loss. Similarly, if the market maker is selling an asset, they must source the asset, again making them vulnerable to the asset's future price movements. Liquidity risk is the main risk market makers encounter, and the bid-ask spread is the way in which market makers are compensated for this risk. When an investor places a market buy (sell) order, the difference between the ask (bid) and the mid price is a premium (concession) for immediate execution.

While the spread is the result of market makers, it in turn affects investor behavior. Amihud and Mendelson (1986) demonstrate that higher yields are required for stocks with larger spreads and that there is a clientele effect whereby stocks with higher spreads are held by investors with longer holding periods. So while a market maker shows a larger spread to compensate for liquidity risk, in a way, this perpetuates a cycle of wide spreads. In the absence of new information, spreads are self-fulfilling as investors do not wish to pay higher transaction prices frequently thereby creating less liquidity. There are a few variations of the bid-ask spread (the quoted spread, the effective spread, and the realized spread) (de Jong and Rindi 2009; Chorida, Roll, and Subrahmanyam 2000), but in essence they all suggest that a large bid-ask spread indicates low liquidity and a small spread suggests high liquidity. The spread measures described below are summarized in Table 2.1.

Quoted Spread

The quoted spread is intended to measure the transactions cost for a round-trip in which one buys (sells) and then sells (buys) the asset at the quoted prices. It is calculated as:

$$QS = p_a - p_b$$

where p_a is the best ask price and p_b is the best bid price. While straightforward, this measure is less indicative of actual round-trip costs both because of the brevity of the asset holding period it assumes and because it does not incorporate any information from trades that have actually been executed. Alternatively, the lack of executed trade information required for computation allows for more frequent measurements, potentially making this measure more relevant in the context of illiquid stocks where trade information is infrequent.

Logarithmic Quoted Spread

Hamao and Hasbrouck (1995) find the distribution of the logarithmic quoted spread to be closer to a normal distribution, so this variation of the quoted spread can be used:

$$\text{Log}QS = \ln(p_a - p_b)$$

However, due to the discretized values of prices and therefore the bid-ask spread, depending on the range of liquidity being observed, it may not always be useful to try to extrapolate a normal distribution from the quoted spread.

Harris (1994) suggested that the minimum price variation can be binding for low price stocks and some frequently traded stocks. This suggests that when the spread widens, it may not widen enough to be reflected through the minimum tick size, thereby limiting the distributional properties of the measure. Because our dataset contains some of the most liquid stocks in the U.S. market such as Bank of America (BAC) and Citigroup (C), we found the quoted spread to consistently be \$0.01 with very little variation even as other measures indicated changes in general liquid-

ity. Trying to fit a normal distribution to these quoted spreads would not only be ineffective but also misleading.

Proportional Quoted Spread

Because the quoted spread uses absolute prices, it is a measure that cannot be compared across stocks of different price ranges. To compare quoted spreads across stocks, one can use the proportional quoted spread. It is calculated as the quoted spread as a percentage of the mid price:

$$PQS = \frac{p_a - p_b}{p_m} = \frac{2(p_a - p_b)}{p_a + p_b}$$

where p_m indicates the mid price. The proportional quoted spread is ubiquitously used in literature as a benchmark liquidity measure.

Relative Quoted Spread

Because the mid is only an estimate of the asset's true value, one might want a spread measure that takes into account realized trades rather than quotes only. The relative quoted spread is identical to the proportional quoted spread, with the exception that the quoted spread is expressed as a percentage of the trade price instead of the mid:

$$RQS = \frac{p_a - p_b}{p_t}$$

where p_t indicates the executed trade price. Because this version of the quoted spread uses executed trade prices, it contains information content of market movement that the other versions of the quoted spreads do not. For frequently traded stocks, this additional information might be more pertinent than for infrequently traded stocks as the long periods of time in between trades might add noise rather than information to this measure.

Effective Spread

As opposed to measuring potential trading costs, the effective spread is a measure of realized trading costs:

$$ES = 2|p_t - p_m|$$

The effective spread and its variations are multiplied by two for consistency with the transaction costs measured by the quoted spread. Additionally, by taking the absolute value of the difference, the measure does not take into account trade direction. The effective spread is a useful measure because it provides an indication of whether trades are occurring at, within, or outside the quoted prices.

Within quote trading is a widely documented practice that is common across asset classes. This is particularly the case in foreign exchange (FX) and bond markets. Because of these markets' lack of centralized exchanges, quoted prices are often merely a starting point for price negotiations.

Because equities markets do have centralized exchanges, the mechanics for trading inside the quoted prices are slightly different for the stock market. For instance, on the New York Stock Exchange (NYSE), market orders may execute at prices within the quotes when the specialist (the NYSE's designated dealer) or a floor broker elects to improve on the quote (Bessembinder and Venkataraman, 2009). Alternatively, many electronic exchanges allow for hidden liquidity, which suggests that limit orders within the quoted spread may exist on the book.

Trading which occurs outside the quoted spread is also common, particularly if the proposed transaction exceeds the quoted depth of shares available to trade. When the transaction size is greater than the quoted depth, the remaining portion of the order may be executed at a different price (Chorida et al., 2000).

So while trading which occurs outside the quotes can happen, within quote trading is still more common than outside quote trading. Chorida et al. (2000) found that the effective spread is somewhat smaller than the quoted spread in an empirical study of the behavior of these different measures on NYSE stocks during the 1992 trading year. Many subsequent studies spanning different time periods, asset classes, and regions have also measured effective spreads to be smaller than quoted spreads.

Proportional Effective Spread

As with the quoted spread, in order to make the effective spread comparable across stocks, it can be computed as a percentage of the mid:

$$PES = \frac{2|p_t - p_m|}{p_m}$$

Relative Effective Spread

The effective spread is also often calculated as a percentage of the trade price:

$$RES = \frac{2|p_t - p_m|}{p_t}$$

Like the relative quoted spread, the relative effective spread does contain directional information since it includes the trade price.

Realized Spread

While the quoted spread represents liquidity conditions for potential trades and the effective spread measures actual trading costs, the realized spread represents the price impact of a trade. Roll (1972) developed a market micro-structure theory that separates the non-informational (inventory and order processing) and informational (adverse selection) components of trading costs. His approach rests on a basic insight: non-informational transaction costs should result only in a temporary deviation of price from value, evidenced by a price reversal after the trade. Meanwhile after informed trading occurs, the price does not mean-revert because it is associated with a permanent price movement.

Due to these adverse price movements, market makers earn less than the effective spreads assuming they hold onto the asset in their inventory for even a brief amount of time. The realized spread is a measure of the realized market making revenue; it can be thought of as the effective spread minus the price impact of the initial trade after a fixed amount of time n :

$$RS = 2|p_t - p_{m_{t+n}}|$$

where p_t is the executed trade price at time t and $p_{m_{t+n}}$ is the mid price at time

$t + n$. The difference is multiplied by two for consistency with the transaction costs measured by other versions of the spread. Additionally, the measure does not take into account trade direction.

The period of time for which price impact is measured, n , is not standard across literature. While five minutes, thirty minutes, and 24 hours are common choices in literature, Hasbrouck (1991) has shown using vector autoregression models of signed trades and bids that it can take many trades for the cumulative price impact to ultimately be realized and that the appropriate window varies depending on the individual stocks.

As with previous spread measures, there are variations of the realized spread in order to compare across stocks. The proportional (relative) realized spread is the realized spread as a percentage of the mid (trade) price at time t .

Proportional Realized Spread

The proportional realized spread is defined as:

$$PRS = \frac{2|p_t - p_{m_{t+n}}|}{p_{m_t}}$$

where p_{m_t} is the mid price at the time of the original trade t .

Relative Realized Spread

Similarly, the relative realized spread is calculated as:

$$RRS = \frac{2|p_t - p_{m_{t+n}}|}{p_t}$$

As mentioned earlier, since the trade price is included in relative spread measures, they include directional information that the other variations of the spread measures do not.

Summary

While the three variations of spread measurements can be thought of as the quoted spread, the effective spread, and the realized spread, there are different versions of these, the most appropriate of which depends on the context. In short, the logarithmic spread can be useful because its distribution is closer to the normal distribution than

Spread Measure	Definition
QS	$p_a - p_b$
LoQS	$\ln(p_a - p_b)$
PQS	$(p_a - p_b)/p_m$
RQS	$(p_a - p_b)/p_t$
ES	$2 p_t - p_m $
PES	$2 p_t - p_m /p_m$
RES	$2 p_t - p_m /p_t$
RS	$2 p_t - p_{m_{t+n}} $
PRS	$2 p_t - p_{m_{t+n}} /p_{m_t}$
RRS	$2 p_t - p_{m_{t+n}} /p_t$

Table 2.1: A brief summary of the spread measures discussed where p is the asset price, a is the ask side, b is the bid side, m is the mid, t indicates an executed trade, and n is an amount of time over which price impact is measured.

that of the quoted spread. Alternatively, the proportional spreads are helpful for comparing spreads across different stocks. The relative spread takes into account market movement by using the trade price instead of the mid price in the denominator of the measure.

However, although these spreads are useful for analysis, they are often difficult to observe in practice. To calculate any of these measures, transaction prices and quote data are needed. These data are not always available and, furthermore, are not always reliable. Because of differences in reporting quote and trade times and exchange conventions, estimating the spread using transaction prices instead of direct calculation from quote data often produces a more representative measure of liquidity. Additionally, the spreads do not always apply to large trades, as the bids and offers are for specific quantities of shares which are relatively small.

Roll (1984) developed a model for estimating spread by computing the autocovariance of transaction prices; however, this model came with strict, unrealistic assumptions, namely that trading does not affect the mid of bid-ask quotes. Even more constraining, Roll's spread cannot be calculated when prices show positive serial correlation. Because it is common practice for traders to split up large orders into a series of smaller trades, thereby creating serial dependence between buy and sell orders, this last condition makes Roll's model simply inapplicable. While extensions

were made to Roll's model, ultimately, intraday data is still needed for the estimation, and the model does not account for positive serial correlation of prices.

2.2 Volume and Depth Measures

While bid-ask spreads are measures of liquidity specific to market makers, volume and depth measures are perhaps the most natural measure of liquidity. The most fundamental of these is of course trading volume. The trading volume in a certain time interval T is:

$$V = \sum_{i=1}^{N_T} v_i$$

where N_T denotes the number of trades in the time interval T and v_i is the number of shares in trade i . A larger trading volume indicates greater liquidity.

In order to scale volume traded across stocks, a stock's turnover can be used as a volume measurement instead. The turnover is the total dollar value of transactions over the time interval T :

$$T = \sum_{i=1}^{N_T} p_i v_i$$

where p_i is the price of trade i .

As with the quoted spread, there is value in quote information, even if it is not traded upon. The depth captures the total available volume to trade at the best bid and best ask prices. It is defined as:

$$D = q_a + q_b$$

where q_a is the best ask quote size and q_b is the best bid quote size. This depth measure is also often expressed as an average depth in which the sum of the bid and ask quote sizes is divided by two. Chordia, Roll, and Subrahmanyam (2000) found that all measures of spread were positively correlated with each other across time and

negatively correlated with depth. A large depth measure indicates greater liquidity.

Dollar depth is an average depth calculation in dollar terms:

$$\$D = \frac{p_a q_a + p_b q_b}{2}$$

Dollar depth makes different stocks' comparable in the same way turnover does to volume. Chordia et al. (2001) use the dollar depth in a measure they call composite liquidity in which the proportional quoted spread is divided by the dollar depth:

$$CL = \frac{2(p_a - p_b)}{p_m \cdot (p_a q_a + p_b q_b)}$$

Composite liquidity combines the two most important aspects of quoted data (price and size) into one liquidity measure. If the minimum tick size does not restrict the absolute spread, then composite liquidity is not affected by the absolute stock price.

The logarithmic version of volume depth is defined as a sum of the logarithms of the best bid and best ask sizes, given by:

$$LogD = \ln(q_a) + \ln(q_b)$$

As with other calculations, taking the log of each component or a combination of them can result in better distributional properties. However, unlike with spreads, the use of the logarithm for depth is arguably more effective for very liquid stocks than it is for illiquid stocks.

Since the best bid and ask quote sizes do not necessarily move in conjunction with one another, the imbalance between them can also be investigated through:

$$I = |p_a * q_a - p_b * q_b|$$

The imbalance between buy and sell volumes is a factor that is used in many more complex liquidity measures (such as the probability of informed trading (PIN) and its volume-based extension (VPIN) introduced by Easley, Lopez de Prado, and O'Hara (2012)). While classifying traded volume as buys or sells is difficult and often impre-

cise, the dollar difference between the bid dollar size and the ask dollar size can give some insight into the buy-sell imbalance.

Ranaldo (2000) extended this concept of buy and ask quote size imbalance to develop the order ratio:

$$OR = \frac{|q_a - q_b|}{p_t \cdot q_t}$$

The order ratio is the ratio of the depth imbalance over turnover. For any given large turnover, if the order imbalance remains small, then liquidity is greater. Thus a large (small) order ratio indicates low (high) liquidity.

2.3 Trade Count Measures

Previous work by Jones, Kaul, and Lipson (1994) suggests that the number of trades, rather than the dollar volume of trading, is a better indicator of individual firm asymmetric information. In fact, they showed that volume had little impact on volatility once trade frequency had been taken into account. Perhaps this result could be explained by informed traders, who in an attempt to obscure the building of their position, split their orders into a series of smaller trades. In other words, large uninformed traders such as institutions might dominate the determination of dollar volume while informed traders might dominate the determination of the number of transactions. Barclay and Warner (1993) suggest that informed traders do break up their orders and are most active in medium-size trades.

The number of transactions per time unit, N_T , is simply defined as the number of trades in the time interval T :

$$N_T = \sum_{i=1}^{N_T} i$$

The reverse of this measure is the waiting time between subsequent transactions. The flow ratio over time interval T is defined as the ratio of turnover to waiting time:

$$FR = \frac{\sum_{i=1}^{N_T} p_i \cdot q_i}{\frac{1}{N-1} \sum_{i=2}^{N_T} tr_i - tr_{i-1}}$$

where tr_i is the time of transaction i . The flow ratio intends to capture whether the volume being traded is taking place in a large number of small trades or in a small number of large trades. Because a larger number of trades corresponds to greater liquidity, a greater flow ratio indicates greater liquidity.

2.4 Price Change Measures

The final category of liquidity measurements we consider consists of those that involve some sort of change in price per volume or trade count. While some literature and models (such as Easley, Lopez de Prado, and O’Hara’s VPIN) suggest that measuring price change over volume bars instead of time bars provides more information, in order to preserve the time structure of the data, which contains information in itself, we will only consider price change measures over fixed units of time. Because many of the measurements that fall under this categorization are very similar, to reduce dimensional, we consider two: the Martin Index and a measurement we will call “liquidity ratio one.”

The Martin Index looks at the squared absolute change in price between consecutive trades and divides this by turnover:

$$M = \sum_{i=1}^T \frac{(p_i - p_{i-1})^2}{p_i * v_i}$$

This measure scales price changes by the turnover, so that the measure for a trade that has a large price impact is adjusted if it is associated with a large turnover. Alternatively, a small trade that has a large price impact will indicate a greater degree of illiquidity. Because the Martin Index squares the price difference between trade prices, it does not take into account directional information.

Given the debate over whether volume or trade count is more meaningful, a measure that examines the price changes divided by the trade count in a period seems appropriate. Brunner (1996) proposes this liquidity ratio that describes the average price change of a transaction over period T :

$$LR = \frac{\sum_{i=1}^N |\Delta P_i|}{N_T}$$

where ΔP_i is the log return from period $i-1$ to i and N_T is the number of trades during time period T . As this measure indicates the average price change of a transaction, a lower liquidity ratio indicates higher liquidity. If the number of trades during a time period is zero, then the liquidity ratio is undefined, as setting it to zero would indicate very high liquidity because of the nature of the measure.

While the previous measures all convey different aspects of intraday liquidity, intraday data can be hard to attain, store, and analyze given the high frequency nature of today's markets. It is impractical to only consider measures that use intraday data. Amihud's (2002) ILLIQ measure is perhaps the most successful liquidity metric that does not require intraday data. Although the measure can be calculated using intraday, it can be (and more often is) calculated using daily data.

The ILLIQ measure aims to compute the price impact of trading. This is accomplished by calculating the daily return on a day, i , divided by the dollar net order flow, the number of shares sold subtracted from the number of shares bought on day i multiplied by the closing price on day i . An average of this ratio is taken over several days to calculate an illiquidity measure for those days. If an exchange does not publish the net order flow, the gross order flow, the total number of buy and sell orders, V_i , is frequently substituted for the net order flow. This gives an ILLIQ measure that looks as such

$$ILLIQ = \frac{1}{N} \sum_{i=1}^N \frac{|\Delta P_i|}{V_i}$$

Under this measure, a low ILLIQ corresponds to small change in price for a given trading volume and therefore indicates high liquidity. Alternatively, a high ILLIQ corresponds to a large change in price for a given trading volume and suggests low liquidity. To compare this measure across stocks, we can think of the measure in terms of dollar flow

$$ILLIQ = \frac{1}{N} \sum_{i=1}^N \frac{|\Delta P_i|}{T_i}$$

where dollar flow T_i on day i is calculated in the same way as turnover is, using the closing price of the stock on day i or the closing mid on day i if available. If ILLIQ is calculated using gross order flow instead of net order flow, ILLIQ will overestimate liquidity, as the gross order flow will always be greater than or equal to the net order flow.

Given its ease of calculation, the ILLIQ measure is particularly useful in demonstrating a basic relationship between liquidity and asset pricing. In a 2006 study, Hasbrouck investigated the correlation between a variety of liquidity measures and found that out of all of the daily liquidity measures, ILLIQ was the most strongly correlated with the effective spread. As an alternative, Hasbrouck (2003) had proposed a Gibbs estimate that is based on a stock's daily closing prices, a modification of Roll's (1984) model of price dynamics. Looking at U.S. Equity trade and quote data that spans 1993-2005, and comprises roughly 300 firms per year (approximately 3,900 firm-years), the ILLIQ measure was most strongly correlated with the effective spread, with the CRSP/Gibbs effective cost estimate being second. Given that the effective spread is generally accepted to be the most accurate liquidity measure estimated from intraday data, ILLIQ is a simpler, attractive alternative.

Chapter 3

Transactions Data for the Financial Sector ETF and Two Component Stocks

The data set we consider for our liquidity modeling is the Financial Select Sector SPDR Fund, which we will often refer to as the Financial Sector ETF (XLF), and two of its component stocks: Bank of America (BAC) and Citigroup (C). Bank of America was selected as one of the component stocks for which to model liquidity, in part, because it was the most liquid ¹ of all of the component stocks that comprise the Financial Sector ETF. However, because Bank of America is so widely traded, at times, its liquidity measures, particularly ratios, can appear unusual, mostly due to the combination of its large volume traded and the extremely small spreads at which it trades. This combination can often lead to very large or very small measures, making some of its results difficult to manage to a certain degree. For this reason, we also include Citigroup stock in our analysis; while it is still a highly liquid stock, we believe it presents a more “normal” model of liquidity for a stock in the U.S. financial sector.

¹Here we define liquidity by a simple heuristic - the the largest average daily volume traded over the 2018-2019 calendar year.

The data for each security consists of trade and quote level data that spans 32 trading days from February 4, 2019 to March 20, 2019 ². For quotes, the best bid and ask quotes, the exchanges of those quotes, the quote sizes, and the date and time of the quotes are included. For trades, the executed trade price, volume of the trade, exchange, and date and time are also included. The accuracy of the times reported for each trade and quote is to the nearest second.

Because the arrival of quote and trade information is irregularly spaced, we overlay five-minute bars onto the data to improve mathematical tractability for measurements. Five-minute bars were chosen as to allow time for liquidity trends to develop for measures which require a time interval (such as trade count, order ratio, ILLIQ, etc.) but also brief enough to accurately reflect more short-term measures such as those related to the spread. Additionally, there is precedent for using five-minute bars as demonstrated by Cao, Hansch, and Beardsley (2004) and Andersen & Bollerslev (1997). The bars include trades that are greater than or equal to the start time of the bar and less than the end time of the bar.

Although the data set has information that extends beyond U.S. market hours, we limit our analysis to transactions during U.S. market hours from 9:30:00 a.m. to 4:00:00 p.m. EST. Therefore the first time bar in the morning extends from 9:30 a.m. until 9:35 a.m. and includes trades occurring at or after 9:30 a.m. and before 9:35 a.m. Since the market closes at 4:00 p.m., there a total of 78 time bars in each day. With 32 trading days of data, each security in the dataset contains 2,496 observations.

Since some of the liquidity measures considered (all quoted and effective spread measures, depth, dollar depth, composite liquidity, log depth, the buy and sell quoted imbalance, and order ratio) are calculated for a single point in time and not defined for calculation over an interval of time, these measures must be aggregated in some way to create 2,496 five-minute bar observations from the two to three million lines of trade and quote data that exist in the original structure of the data.

For this reason, when the five-minute grid is superimposed onto the data, the means of the individual spread, composite liquidity, and order ratio observations

²These data were kindly provided by Kempthorne Analytics.

in each bar are taken to produce one observation for each of the 2,496 five-minute bars. The number of observations that is averaged to create each five-minute bar is variable as it depends on the underlying flow of trade and quote data. Thus the means represented for these measures later in this chapter are averages of five-minute averages.

Due to the nature of the depth, dollar depth, log depth, and the buy and sell quoted size imbalance, these measures are summed, not averaged, over the course of each five-minute time bar to create 2,496 five-minute bar observations. As with the spread and other composite measures, the number of observations that are summed to create each five-minute bar for these depth measures is variable and depends on the underlying quote flow during the specific five-minute bar.

3.1 Financial Sector Exchange-Traded Fund (XLF)

A spider (SPDR) is an exchange-traded fund (ETF) that aims to track the Standard & Poors 500 (S&P) Composite Stock Index. While SPDR is an acronym for S&P Depository Receipts, the term more colloquially refers to any ETF that tracks the S&P. The Financial Select Sector SPDR Fund is a SPDR ETF that attempts to simulate the price and yield performance of the Financial Select Sector Index, an index composed of a subset of companies from the S&P that are classified as financials. The financial industry categorization entails companies from diversified financial services, insurance, commercial banks, capital markets, real estate investment trusts, thrift and mortgage finance, consumer finance, and real estate management and development. The 20 largest constituents by market share in the ETF as of March 22, 2019 are in Table 3.1.

The Financial Select Sector SPDR Fund is a physical ETF with the fund investing substantially all, at least 95%, of its total assets in the securities comprising the index. This is notable as some ETFs synthetically replicate the index and as such their

Stock	Ticker
AIG	AIG
American Express	AXP
Bank of America	BAC
Bank of New York Mellon	BK
BlackRock	BLK
Berkshire Hathaway	BRK.B
Citigroup	C
Chubb	CB
CME Group	CME
Capital One	COF
Goldman Sachs	GS
JPMorgan Chase	JPM
Metlife	MET
Marsh & McLennan	MMC
Morgan Stanley	MS
PNC Financial Services Group	PNC
Prudential Financial	PRU
Charles Schwab	SCHW
U.S. Bancorp	USB
Wells Fargo	WFC

Table 3.1: The largest constituents that make up the Financial Select Sector SPDR Fund as of March 22, 2019. Only Bank of America and Citigroup are included in this analysis.

Name	Ticker	Market Value Weight (%)
Berkshire Hathaway Inc B	BRK.B	12.54%
JPMorgan Chase & Co	JPM	11.12%
Bank of America Corporation	BAC	8.50%
Wells Fargo & Co	WFC	6.77%
Citigroup Inc	C	5.00%
US Bancorp	USB	2.52%
American Express Co	AXP	2.42%
Goldman Sachs Group Inc	GS	2.18%
CME Group Inc Class A	CME	2.08%
Chubb Ltd	CB	1.98%

Table 3.2: The top 10 holdings by market value of XLF as of March 22, 2019.

components and the ETF itself might not have liquidity models and measures that are as related and interdependent. As of March 22, 2019, the portfolio is comprised of 68 holdings and the top 10 holdings by market share are listed in Table 3.2. The returns of the ETF track the underlying index fairly well with the average tracking

difference falling in the 93% percentile. The Financial Select Sector SPDR Fund is the most liquid of the financial sector ETFs with average daily volumes over a variety of time periods consistently greater than 50,000,000 shares. For comparison, the other well-known financial ETFs, the Vanguard Financial ETF (VFH) and the S&P Regional Banking SPDR ETF (KRE), have volumes that are orders of magnitudes smaller.

3.2 Exchange-Traded Fund (ETF) Market Mechanics

Modeling the liquidity of a sector ETF and two of its components can be interesting because of the mechanics of the ETF market and the implications of those mechanics on the liquidity of ETFs component stocks. Although ETFs were first launched in 1993, generally speaking, they were not widely used until after the financial crisis. Because ETFs are funds that aim to track an index or basket of assets, in some sense, they are similar to mutual funds. However, while the creation and redemption of mutual fund shares can occur only at market-close, ETFs trade on exchanges throughout the day. Because of this, unlike for a mutual fund whose share price is determined by its net asset value (NAV) at the end of each day, the price of the ETF fluctuates intraday is determined by market participants transacting in the secondary market.

ETFs are managed by a fund adviser (the Financial Sector ETF's is State Street) and ETF shares can only be created or redeemed by an authorized participants (AP). Usually these APs are large broker dealers with the ability to source assets more cheaply than institutional and retail market participants. To create shares of an ETF, an AP buys the correct proportions of the underlying constituents that make up the ETF in the market. The AP then sells this physical full replication of ETF shares to the ETF provider at the NAV of the basket, not the price of the ETF itself. In exchange, the ETF provider gives ETF shares to the AP, which the AP can then

sell in the secondary market.

This process can also work in reverse to accommodate the redemption of ETF shares. The AP can remove ETF shares from the market by purchasing enough of those shares to form a creation unit. After delivering those shares to the ETF issuer, the AP receives the same value in the underlying securities of the fund. Since the ETF is only a wrapper for this physical basket of securities, the ETF trades separately than the securities do and due to fluctuations in supply and demand, the price of the ETF can (and does) differ from the fund's underlying NAV. For this reason, the creation and redemption process of ETFs is crucial.

When an ETF's price deviates from the NAV, APs use the creation and redemption mechanism to arbitrage the pricing difference, consequently bringing the price back in line with the NAV. If an ETF is trading above the NAV (at a premium), then an AP will buy the underlying ETF components in the market and create ETF shares, thereby receiving the higher-priced ETF in exchange for the lower-valued NAV. Increasing the supply of ETF shares naturally moves the price lower. Alternatively, if the ETF is trading below the NAV (at a discount), an AP will redeem the ETF and sell its constituents in the market to earn profit.

While this arbitrage mechanism works fairly effectively under normal market conditions, the relationship between an ETF and its NAV can deteriorate during times of market stress when the market maker's inventory concerns outweigh the potential benefit from arbitraging the difference between an ETF and its underlying value. This relationship between a sector ETF, its underlying components, and overall sector liquidity is what makes a selection of a sector ETF and a few of its larger components particularly interesting for liquidity modeling.

3.3 ETF Asset Valuation

Although arbitrage is a mechanism the ETF market relies on for efficient ETF pricing, it inextricably ties all of the ETF constituents to each other, arguably creating less efficient pricing for the components themselves. This is particularly the case where

investors are looking for a way to quickly express a general macroeconomic view (for example on a specific sector, asset class, or region), where the intention is to receive exposure to a market theme, not a specific firm. ETFs offer on-demand liquidity and their volumes often peak around events, such as surprise policy decisions or economic data releases where liquidity might be difficult to source otherwise.

Perhaps one of the most dramatic examples of this occurred during the beginning of the financial crisis. On September 19, 2008, due to unprecedented market turmoil, particularly in the financial sector, the Securities and Exchange Commission (SEC) halted short selling of financial-related stocks, specifically those of the bulge bracket banks. In turn, the volume traded on the XLF ETF surged, where market participants, not able to short sell the underlying components, transitioned to short-selling the ETF. In fact, the day when the SEC announced its ban on short selling financials resulted in the largest daily volume traded for the XLF since its inception.

However, the XLF did not only contain bulge bracket banks and the price and liquidity of some of its stock components from sub-sectors not yet affected by the financial crisis were inevitably affected by the ETF transactions. For instance, studies have highlighted concerns related to the pricing and trading of these instruments, including the more rapid transmission of liquidity shocks, higher return correlations among stocks held by same ETFs (Da and Shive 2013, Sullivan and Xiong 2012), greater systemic risk (Ramaswamy 2011), and elevated intraday return volatility both for the component stocks and for the entire market (Ben-David et al. 2014, Broman 2013, Krause et al. 2013), particularly in times of market stress (Wurgler 2010). These different liquidity dynamics of ETFs and their stock components make for particularly interesting liquidity modeling problems with the potential for interdependencies between the securities to actually aid in the modeling process.

Chapter 4

Summary Statistics of Liquidity Measures

In this chapter, we apply the liquidity measures described in Chapter 2 to the aggregated five-minute time bars of trade and quote data for Bank of America, Citigroup, and the Financial Sector ETF. Because there are 78 five-minute time bars in each trading day, and the data spans 32 days, there are 2,496 observations for each liquidity measure. For each measure, the mean, median, maximum, minimum, standard deviation, coefficient of variation, skewness, and kurtosis are calculated on the 2,496 five-minute bar values. In order to understand these measures in an intuitive sense, these measures are initially left non-normalized. Non-normalized measures in their original units provide market context, which is useful to have before moving to analysis in later chapters that involves increasing levels of abstraction.

Since the measures are not normalized, the statistics for many of the measures are inevitably sensitive to the scale of the specific measure and the absolute price of the stock. For this reason, we include the coefficient of variation, skewness, and kurtosis in the statistics summary. These statistics are unitless and therefore comparable across different measures and different stocks. The coefficient of variation is the ratio of the standard deviation to the mean, acting as a relative measure of the variability in the measure. The skewness and kurtosis are indicators of the degree to which a measure

does not follow a normal distribution, due to asymmetry and/or heavy tails; the skewness and kurtosis of a normal distribution are 0 and 3 respectively. Because the aggregate data for each security is mixture over 32 days as well as intraday periods, the distributions of the measures are expected to be complex mixtures. As such, they are likely to have significant skewness and/or kurtosis depending on the analysis period.

To understand the how the measures vary over the analysis period, we show how the daily averages, standard deviations, and coefficients of variation change over the entire time period. Additionally, through the summary statistics and the plots the measures over time, it becomes clear that some of these measures are highly correlated. Specifically, across all three securities, the following groups of measures appear to be effectively equivalent: (1) the quoted spread, the proportional quoted spread, the relative quoted spread, and the relative spread of log prices; (2) the effective spread, the proportional effective spread, and the relative effective spread; (4) the realized spread, the proportional realized spread, and the relative realized spread; (5) volume and turnover; and (6) depth and dollar depth.

Finally, to examine how these measures behave on an intraday basis, rather than over the course of the whole time period, each day's measure values for the 78 five-minute bars are converted to Z-scores. The Z-scores adjust a measure's values over the 32 days by normalizing each day's values to have mean zero and unit standard deviation. As a result, day-to-day variation is removed from the measures, which enables the comparison of intraday variation across days. 15-minute bar averages across the 32 days are then taken to examine liquidity statistics for intraday variation. Z-scores are averaged over 15-minute bars, instead of five-minute bars, to increase the intraday sample sizes per period from 32 (the number of days) to 96.

From this transformation, we are able to see systematic intraday liquidity patterns in all securities that seem to persist across days regardless of the level of liquidity of a specific day. The most noticeable of these patterns is the U-shape that appears in the volume and trade count measures, where over the course of the trading day, the measures indicate high liquidity at times near the open and close of trading and low

liquidity in the middle of the day.

This U-shaped pattern is one that has been widely researched with Wood et al. (1985) and Harris (1986) noting the pattern's existence in stock return volatility in as early as 1985. And while the U-shape was first noticed in return volatility, it has since been shown that the pattern of intraday return volatility is highly correlated with the intraday variation of trading volume and bid-ask spreads (Andersen and Bollerslev 1997). Although we are examining securities from the equity market, the U-shape is ubiquitous across asset classes with its existence equally as pronounced in foreign exchange and fixed income markets as demonstrated by Müller et al. (1990) and Baillie and Bollerslev (1991).

4.1 Liquidity Measures for Bank of America

Tables 4.1 through 4.4 show the statistic summaries for Bank of America's liquidity measures. As in Chapter 2, the measures are categorized into spread measures, volume and depth measures, trade count measures, and price change measures; each category's statistics are presented individually. These statistics are taken across the five-minute bars of the non-normalized, absolute measures, and encompass 2,483 observations. While there are 2,946 time bars in the time period, the last 13 five-minute time bars, or equivalently 65 minutes, of the last day's worth of data were omitted from the original data set thus resulting in a sample size of 2,483. Instead of interpolating the undefined values, we choose to simply not include observations for each liquidity measure of Bank of America during the last 13 five-minute periods on March 20, 2019, the final day of the analysis.

The summary statistics for Bank of America's spread measures are displayed in Table 4.1. The average quoted spread of Bank of America is \$0.01 while its average effective spread is slightly smaller at \$0.005. The average effective spread being less than the average quoted spread indicates that within quote trading consistently occurs. Beyond just within quote trading, because the average effective spread is less than the minimum tick size of \$0.01, it suggests that a large number of trades

	QS	PQS	RQS	RSLogP	ES
Mean	1.00e-02	3.45e-04	3.45e-04	3.45e-04	4.48e-03
Median	1.00e-02	3.45e-04	3.45e-04	3.45e-04	4.50e-03
Max	1.10e-02	3.85e-04	3.85e-04	3.85e-04	5.23e-02
Min	1.00e-02	3.32e-04	3.32e-04	3.32e-04	0.00e+00
Std Dev	5.59e-05	5.94e-06	5.94e-06	5.94e-06	2.04e-03
Coeff Var	0.01	0.02	0.02	0.02	0.46
Skewness	9.57	0.68	0.68	0.68	5.32
Kurtosis	121.49	1.26	1.26	1.26	120.53
	PES	RES	RS	PRS	RRS
Mean	1.55e-04	1.55e-04	4.80e-02	1.66e-03	1.66e-03
Median	1.55e-04	1.55e-04	3.00e-02	1.05e-03	1.05e-03
Max	1.77e-03	1.79e-03	3.90e-01	1.36e-02	1.36e-02
Min	0.00e+00	0.00e+00	0.00e+00	0.00e+00	0.00e+00
Std Dev	7.01e-05	7.04e-05	4.79e-02	1.66e-03	1.66e-03
Coeff Var	0.45	0.46	1.0	1.0	1.0
Skewness	5.08	5.26	2.38	2.4	2.4
Kurtosis	113.35	118.69	8.67	8.84	8.84

Table 4.1: Summary statistics of Bank of America’s spread measures; these encompass 2,483 five-minute bar observations for each measure.

were executed at the mid quote, which is considered to be the fair value price. The maximum and minimum average quoted spreads across five-minute time bars are both \$0.01. Because the maximum quoted spread here is a maximum of five-minute time bars that contain the average quoted spread during those five minutes, the \$0.01 maximum does not indicate that the quoted spread never exceeded \$0.01, but rather that the average quoted spread in five minutes never did.

While the average quoted spread is near \$0.01, the average realized spread is \$0.048, with the median realized spread equaling \$0.03. Because the realized spread is a measure of where the quote mid price is five minutes after a trade occurs multiplied by two, we can consider the average move in price to be around \$0.02. Since the quoted and effective spreads are smaller than \$0.02, which would correspond to profit losses for market makers, it would stand to reason that market makers are keeping this stock in their inventory over a time horizon longer than five minutes.

Figure 4-1 shows the evolution of the daily average, daily standard deviation, and daily coefficient of variation for the quoted, effective, and realized spreads, and from

this, it appears that there is very little variation in the average quoted and effective spread measures. Figure 4-1 displays these values on a scale that demonstrates their variation, but, the scale is so small for the quoted and effective measures, especially in the context of a \$0.01 minimum tick size, that practically speaking, there is not much variation in these average spread measures over the period of analysis.

Harris (1994) commented on the effect of price discretization for very liquid stocks, suggesting that the minimum price variation can in fact be binding for low price stocks and some frequently traded stocks. He concluded that when the spread widens, it may not widen enough to be reflected through the minimum tick size, limiting the information content in these spread measures. The static nature of Bank of America's quoted and effective spread measures in combination with its large volume traded suggest that the discretization of prices, which in turn discretizes the bid-ask spread, may be limiting the information content in Bank of America's quoted and effective spreads.

In terms of variation, the coefficient of variation for the realized spread is much greater than those for the quoted and effective spreads. From Figure 4-1, it is evident that the realized spread's coefficient of variation spans the largest range, followed by the effective spread, and finally the quoted spread.

All of the spread measures are skewed to the right. Because the quoted spread and effective spread average values are so small and constrained to being greater than or equal to zero by definition in the case of the effective spread, this in turn greatly reduces the possibility of negative skew. This in part also explains the very large kurtosis values of the quoted and effective spread measures as they indicate fat-tailed distributions.

As Bank of America's spread measures are very small, its volume and depth measures are correspondingly quite large. Table 4.2 displays the statistics for these measures. The average number of shares traded in five minutes is 555,000 which translates to an average turnover of \$16,100,000 every five minutes. Bank of America is one of the most widely traded stocks in the world and has the largest average volume out of all of the securities examined in this analysis. Therefore it is not unexpected

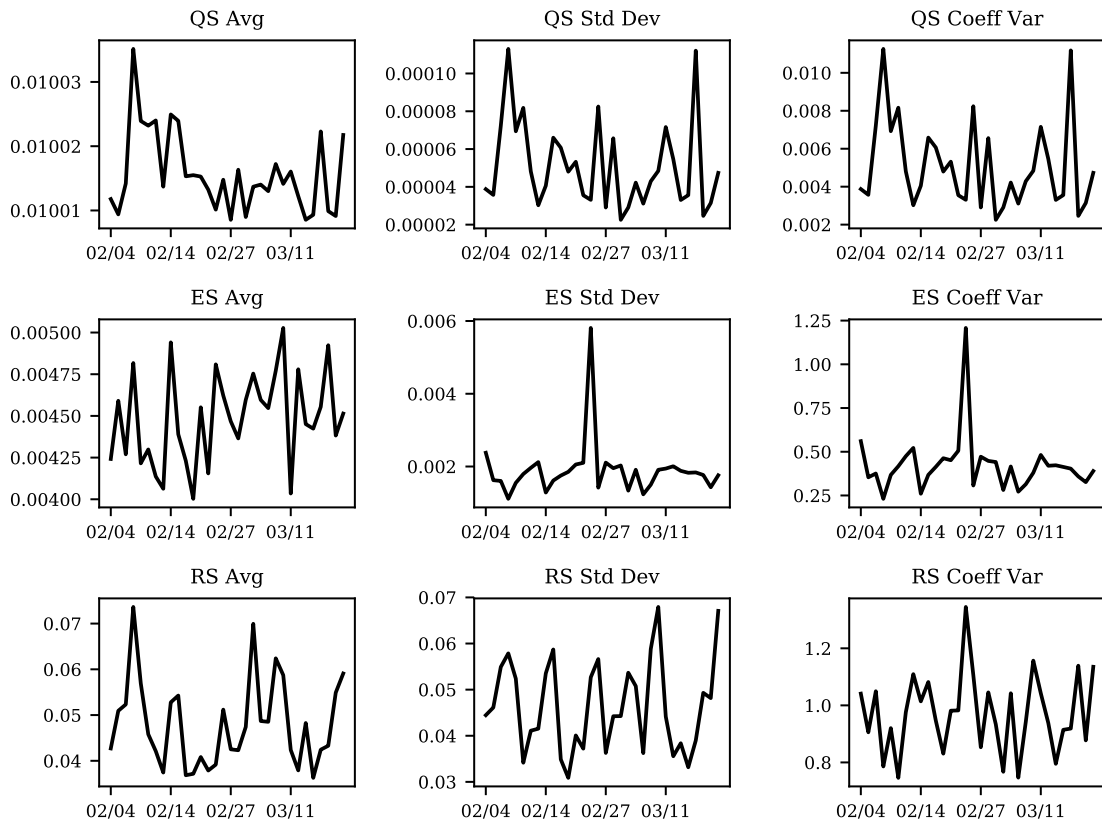


Figure 4-1: The average, standard deviation, and coefficient of variation of Bank of America's spread measures, aggregated across 2,483 five-minute bars.

	V	T	D	\$D
Mean	5.07e+05	1.47e+07	5.98e+05	8.68e+06
Median	4.00e+05	1.16e+07	4.54e+05	6.61e+06
Max	9.03e+06	2.66e+08	1.16e+07	1.69e+08
Min	1.01e+04	2.95e+05	2.60e+03	3.80e+04
Std Dev	4.62e+05	1.35e+07	6.98e+05	1.02e+07
Coeff Var	0.91	0.91	1.17	1.17
Skewness	6.92	7.02	7.94	7.96
Kurtosis	89.25	92.04	86.48	87.15
	CL	LogD	I	OR
Mean	6.88e-08	2.48e+01	2.15e+06	2.72e-01
Median	6.09e-08	2.46e+01	1.36e+06	2.27e-01
Max	8.71e-07	3.11e+01	5.99e+07	2.24e+00
Min	2.02e-08	1.42e+01	1.01e+02	2.66e-02
Std Dev	5.80e-08	1.11e+00	3.39e+06	1.74e-01
Coeff Var	0.84	0.04	1.58	0.64
Skewness	0.0	1.02	7.38	2.53
Kurtosis	0.0	7.48	85.21	13.21

Table 4.2: Summary statistics of Bank of America’s volume and depth measures; these encompass 2,483 five-minute bar observations for each measure.

that the standard deviations for both volume and turnover are very large. Figure 4-2 shows the daily average, standard deviation, and coefficient of variation across all depth measures over the course of the analysis time period.

Composite liquidity, which is the proportional quoted spread over dollar depth, is essentially zero across all statistics. This is because the dollar depth dwarfs the proportional quoted spread. However, while the scale is very small, composite liquidity does seem to demonstrate a fair amount of variation. Its coefficient of variation is 0.84, which is almost as large as the coefficients of variation of volume and turnover exhibit.

The average dollar difference between the best bid and best ask sizes being quoted in five minutes is very large at \$2,150,000, or on average 27.2% of the turnover traded as indicated by the average order ratio. This order ratio is unusually large, as one would expect a liquid stock such as Bank of America to have relatively equal best bid and ask sizes; perhaps this indicates that at times market makers are skewing their bid-ask sizes in reaction to increased or decreased liquidity risk. While market makers

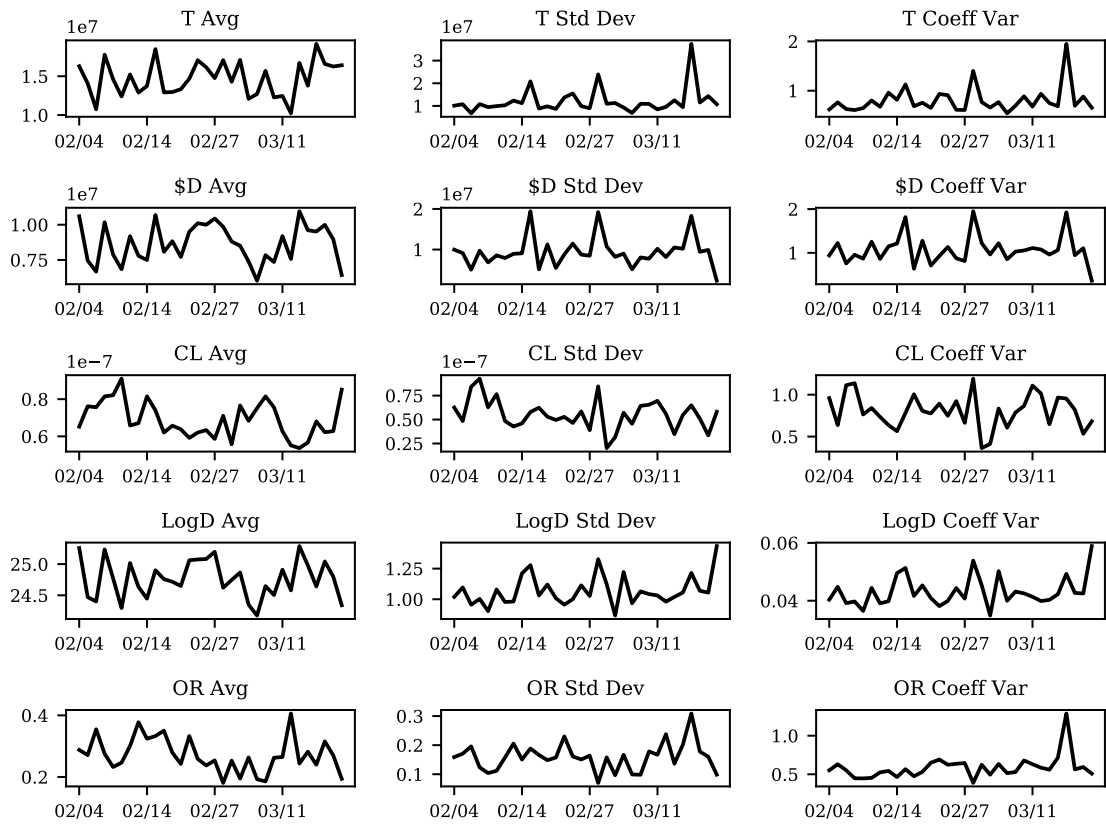


Figure 4-2: The average, standard deviation, and coefficient of variation of Bank of America's volume and depth measures, aggregated across 2,483 five-minute bars.

	N	FR
Mean	1.47e+03	1.20e+08
Median	1.22e+03	5.21e+07
Max	1.52e+04	1.19e+10
Min	8.00e+00	2.36e+06
Std Dev	9.69e+02	3.95e+08
Coeff Var	0.66	3.28
Skewness	4.56	20.33
Kurtosis	42.83	540.52

Table 4.3: Summary statistics of Bank of America’s trade count measures; these encompass 2,483 five-minute bar observations for each measure.

do not seem to show much variation in their spread, they do seem to show quite a bit of variation in their bid-ask sizes. The quoted depth sizes, but not spreads, of Bank of America adjusting for liquidity further supports the notion that when Bank of America’s spread widens, it may not widen enough to be reflected through the minimum tick size; as such, the depth may be an easier means through which market makers can reflect liquidity variations.

Shown in Table 4.3, the average number of trades during each time bar is 1,470 with a minimum of 8 and a maximum of 1,520. The flow ratio, which is the ratio of turnover to average trade waiting time is on average very large at \$120,000,000, and indicates that the volume traded occurs over series of small trades rather than large, infrequent trades. Other than the trade count and flow ratio measures, block trades are not identified or accounted for in any specific way in this analysis. However, it is generally accepted practice to split large orders into a series of smaller ones, the serial autocorrelation of which we capture in our analysis of time series analyses.

Table 4.4 displays the statistics for the price change measures. These measures appear to have incredibly small values across all statistics, most likely due to the fact that all of the volume and depth measures overshadow the small price changes and spread measurements. However, the coefficient of variation, skewness, and kurtosis values for all of these measures are incredibly large. This suggests a very large degree of variability in the measures, expressed through fat-tailed distributions.

Liquidity ratio one, which examines the average price change of a trade, indicates

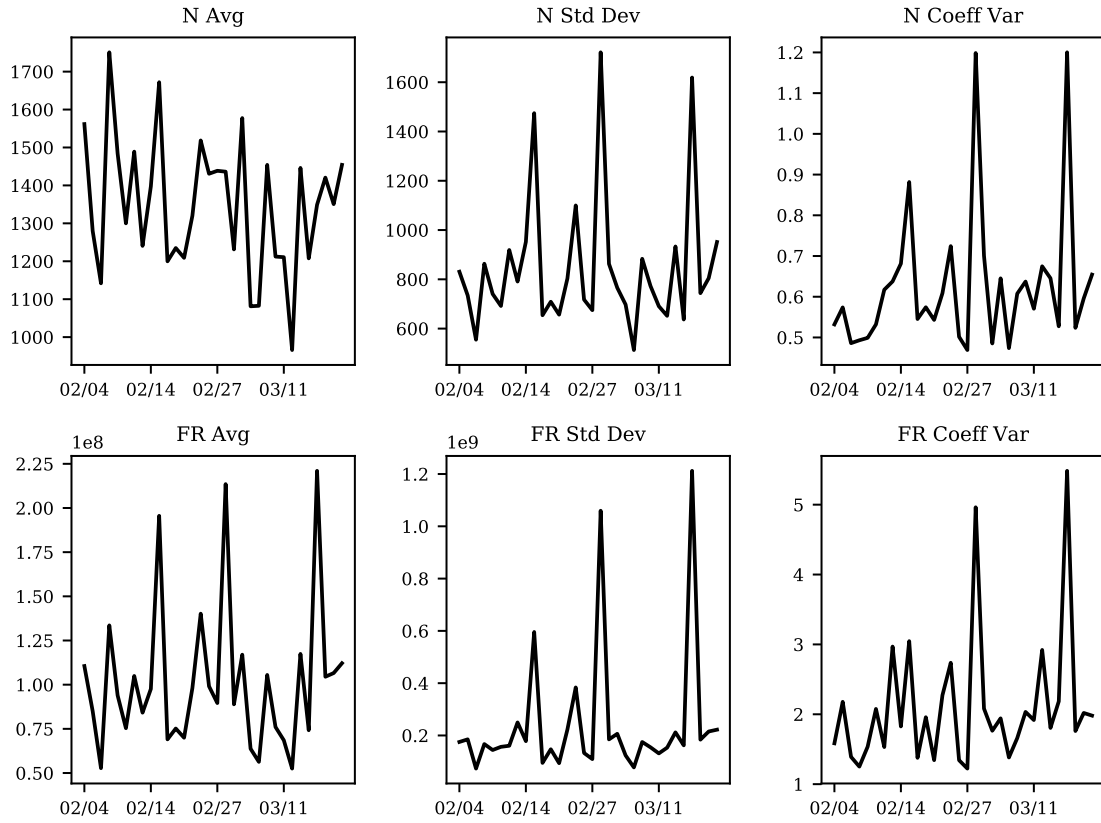


Figure 4-3: The average, standard deviation, and coefficient of variation of Bank of America's trade count measures, aggregated across 2,483 five-minute bars.

	M	LR1	Illiq
Mean	4.86e-04	5.32e-08	6.62e-11
Median	2.70e-05	4.22e-08	5.32e-11
Max	1.42e-01	1.06e-05	6.85e-10
Min	7.28e-07	3.10e-09	0.00e+00
Std Dev	4.19e-03	2.14e-07	5.72e-11
Coeff Var	8.62	4.01	0.86
Skewness	20.9	0.0	0.0
Kurtosis	593.61	2385.65	0.0

Table 4.4: Summary statistics of Bank of America's liquidity measures involving price changes; these encompass 2,483 five-minute bar observations for each measure.

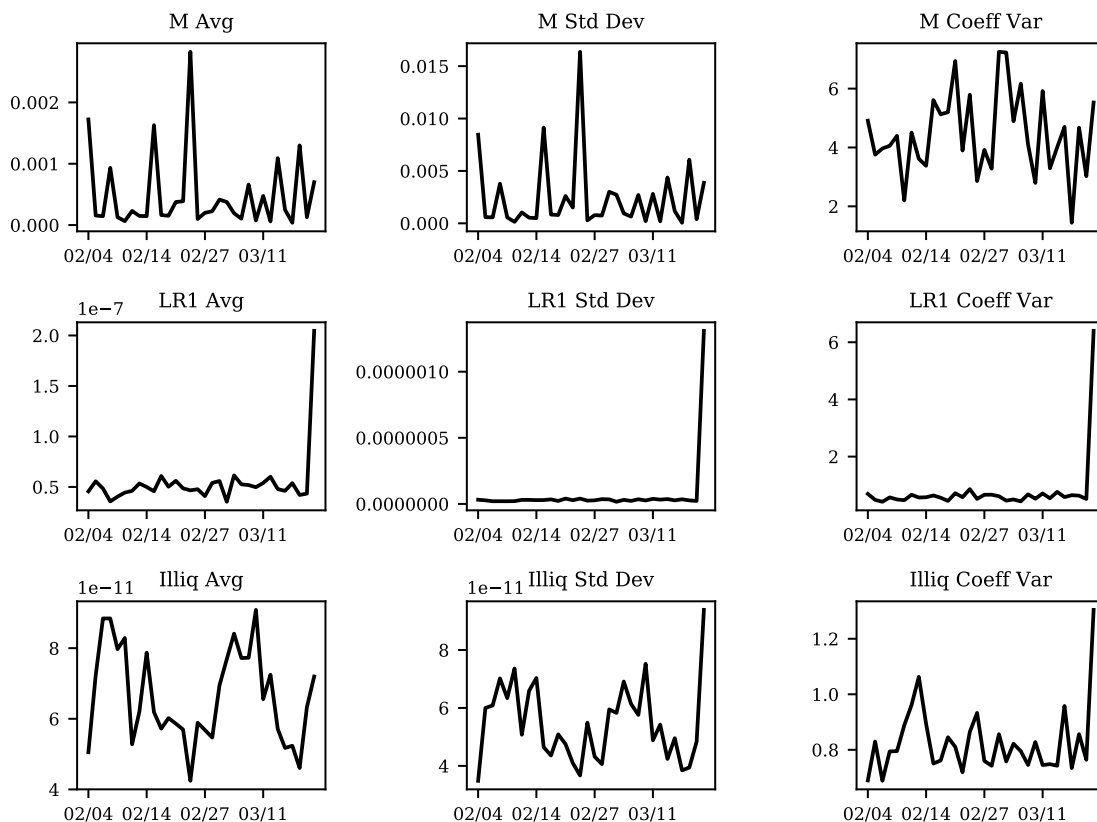


Figure 4-4: The average, standard deviation, and coefficient of variation of Bank of America’s price-change measures, aggregated across 2,483 five-minute bars.

that the average price change of a transaction is effectively zero; however the coefficient of variation is large, and even more so, the kurtosis is incredibly high. Liquidity ratio one and ILLIQ, the five-minute bar return divided by turnover, have zero skew and ILLIQ has a kurtosis of 0.

Although Tables 4.1 through 4.4 provide insight into how Bank of America’s liquidity measures evolve over the 32-day period, these measures are susceptible to variation across days. To investigate the behavior of these liquidity measures in an intraday context, we can examine the Z-scores of each measure; these Z-scores are calculated from 15-minute aggregates and have been standardized across each day. They are displayed in Figure 4-5.

As suggested by previous literature, volume and turnover exhibit a strong U-shape pattern to their intraday liquidity. Trade count, which has shown to be positively

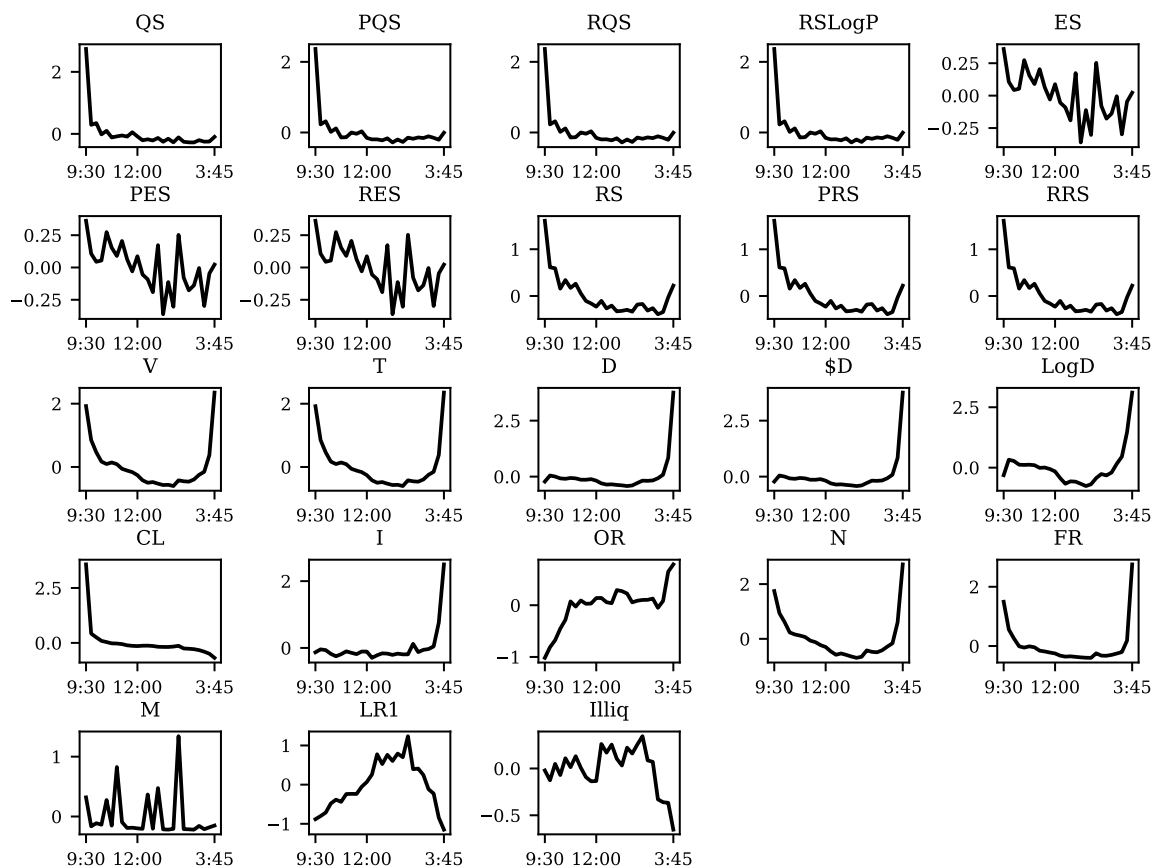


Figure 4-5: The Z-scores of Bank of America’s liquidity measures averaged across 15-minute intervals.

correlated with volume, and flow ratio, a trade-count-related measure, demonstrate this same U-shape pattern. Liquidity ratio one’s inverted U-shape further solidifies the notion of high liquidity near market open and close and low liquidity in the middle of the day because, since it is the average price change of a transaction, a low value indicates greater liquidity.

The spread measures, however, appear to decrease over the course of the trade day, indicating low liquidity in the morning and a consistent level of higher liquidity in the afternoon and near market close. Perhaps, the spreads are greater in the morning, not due to a lack of volume, but rather the possibility of asymmetric information. News often comes out after market close or overnight and economic data is released in the morning, which could potentially explain the greater spreads relative to the volume trading.

4.2 Liquidity Measures for Citigroup

Tables 4.5 through 4.8 include the statistics of Citigroup's liquidity measures on 2,485 aggregated five-minute time bars. As with Bank of America, a time period of information was omitted from the last day's worth of data in the analysis, and these 55 minutes, or 11 five-minute bars, of missing data on the last day of the analysis are not adjusted for in any way.

The statistics of the spread measures are in Table 4.5 and the average, standard deviation, and coefficient of variation of the quoted, effective, and realized spreads over the 32-day analysis period are displayed in Figure 4-6. The average quoted spread for a five-minute bar is \$0.01 and the average effective spread is \$0.009. The average effective spread being less than the minimum tick size indicates that trading at mid occurs frequently. However, given the average effective spread is only slightly less than \$0.01, within quote trading is certainly not occurring the majority of the time. Meanwhile, the realized spread, which measures two times the absolute value of the difference between trade price and the mid quote five-minute later, is \$0.12.

When examining the value of the average spread measures over the course of the entire 32 days, there is some variation, but the ranges which these measures move in is quite small, as indicated by the small magnitude of their coefficients of variation. Over the time period, the spread measures appear to increase and decrease in unison. It appears there were at least two noticeable periods of low liquidity as indicated by the two peaks in daily average values, most clearly seen in the effective and realized spreads. The peaks in spread are much more limited in the average daily spread to the point that this liquidity trend most likely would not have been evident from examining only the quoted spread alone.

However, because all of these spread measures do move in unison, it does suggest a liquidity trend occurred over the time period. While there appears to be a fair amount of noise, or rather uninterpretable movement, in the ratio of the standard deviation to mean for the quoted and realized spread, it constant for the majority of the period for the effective spread. Although the daily average value of the effective

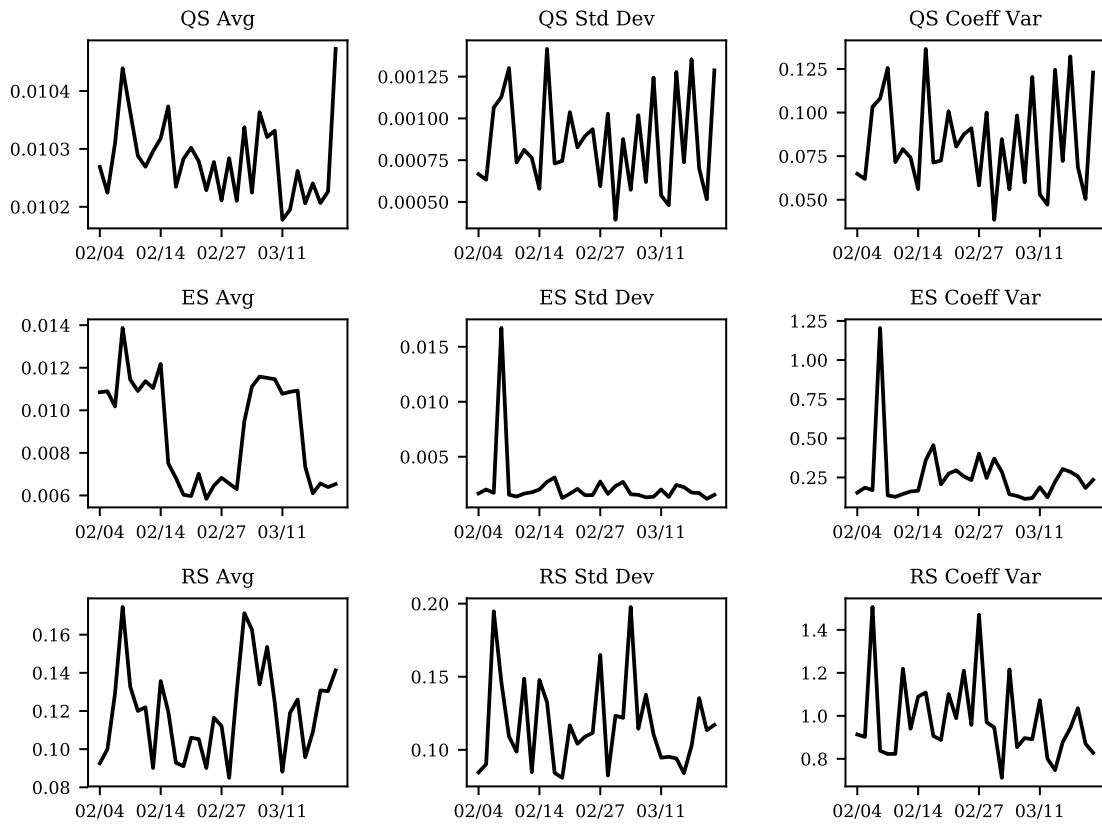


Figure 4-6: The average, standard deviation, and coefficient of variation of Citigroup's spread measures, aggregated across 2,485 five-minute bars.

	QS	PQS	RQS	RSLogP	ES
Mean	1.03e-02	1.62e-04	1.62e-04	1.62e-04	9.03e-03
Median	1.01e-02	1.59e-04	1.59e-04	1.59e-04	9.45e-03
Max	2.25e-02	3.56e-04	3.56e-04	3.56e-04	1.49e-01
Min	1.00e-02	1.50e-04	1.50e-04	1.50e-04	2.12e-03
Std Dev	8.99e-04	1.45e-05	1.45e-05	1.45e-05	4.25e-03
Coeff Var	0.09	0.09	0.09	0.09	0.47
Skewness	9.22	8.61	8.6	8.6	15.39
Kurtosis	93.49	86.38	86.37	86.38	484.23
	PES	RES	RS	PRS	RRS
Mean	1.43e-04	1.43e-04	1.20e-01	1.88e-03	1.88e-03
Median	1.50e-04	1.50e-04	9.00e-02	1.40e-03	1.40e-03
Max	2.38e-03	2.33e-03	1.70e+00	2.66e-02	2.67e-02
Min	3.28e-05	3.28e-05	0.00e+00	0.00e+00	0.00e+00
Std Dev	6.85e-05	6.79e-05	1.22e-01	1.92e-03	1.92e-03
Coeff Var	0.48	0.48	1.02	1.02	1.02
Skewness	14.96	14.5	3.38	3.37	3.38
Kurtosis	465.88	446.61	24.07	23.84	23.96

Table 4.5: Summary statistics of Citigroup’s spread measures; these encompass 2,485 five-minute bar observations for each measure.

spread moves around quite a lot, its standard deviation is constant with the exception of one day of presumably unusually high liquidity.

Citigroup’s spread measures are all significantly skewed to the right and demonstrate very large kurtosis values. This is consistent with the almost zero average quoted and effective spreads. These distributions are heavily skewed, because their average is at their lower bound by definition. To a certain extent, perhaps these incredibly large kurtosis values support the notion that the minimum tick size is restricting the spread on very liquid stocks, such as Bank of America and empirically Citigroup as well.

For volume and depth measures, as indicated in Table 4.6, during Citigroup’s average five-minute bar, 138,000 shares, or rather \$8,780,000 in turnover is traded. Composite liquidity values are very close to zero because the dollar depth in the denominator of the measure is much greater than the proportional quoted spread. The average imbalance dollar amount between quoted bid and ask sizes is \$289,000 and 13.6% of the its turnover. With the exception of composite liquidity which has

	V	T	D	\$D
Mean	1.38e+05	8.78e+06	2.89e+04	9.19e+05
Median	1.06e+05	6.73e+06	1.94e+04	6.19e+05
Max	3.34e+06	2.18e+08	1.37e+06	4.45e+07
Min	3.00e+04	1.93e+06	6.80e+03	2.17e+05
Std Dev	1.33e+05	8.49e+06	5.08e+04	1.62e+06
Coeff Var	0.96	0.97	1.76	1.76
Skewness	10.01	10.34	12.45	12.72
Kurtosis	185.69	197.42	238.79	252.69
	CL	LogD	I	OR
Mean	2.28e-07	1.85e+01	2.89e+05	1.36e-02
Median	1.88e-07	1.83e+01	1.45e+05	1.12e-02
Max	2.61e-06	2.69e+01	2.84e+07	1.18e-01
Min	5.15e-08	1.63e+01	1.97e+02	1.11e-03
Std Dev	1.90e-07	1.15e+00	8.53e+05	9.82e-03
Coeff Var	0.84	0.06	2.96	0.72
Skewness	0.0	2.06	18.99	3.33
Kurtosis	51.46	7.17	526.48	21.15

Table 4.6: Summary statistics of Citigroup’s volume and depth measures; these encompass 2,485 five-minute bar observations for each measure.

zero skew, these measures are all strongly right skewed and have significant fat tails. As with spreads, it appears that unusual volume amounts tend to surprise to the upside rather than the downside.

Figure 4-7 demonstrates the daily means of these volume and depth measures, and while they are clearly fluctuating over time, the liquidity trends are not quite as steady and interpretable as they were for the spread measurements. However, one noticeable feature is that there are distinguished peaks (or minimums depending on the direction of the measure that indicates high liquidity) in all of the average values of the measures at the end of the period on March 15, 2019. However, the standard deviations and coefficients of variations around this time, only show sharp increases for turnover, dollar depth, and order ratio.

Because this peak in value is not an unusual dispersion for composite liquidity and log depth, as indicated by the lack of change in their standard deviations and coefficients of variation, perhaps this suggests that composite liquidity and log depth may capture a different dimension of Citigroup’s depth and volume liquidity than the

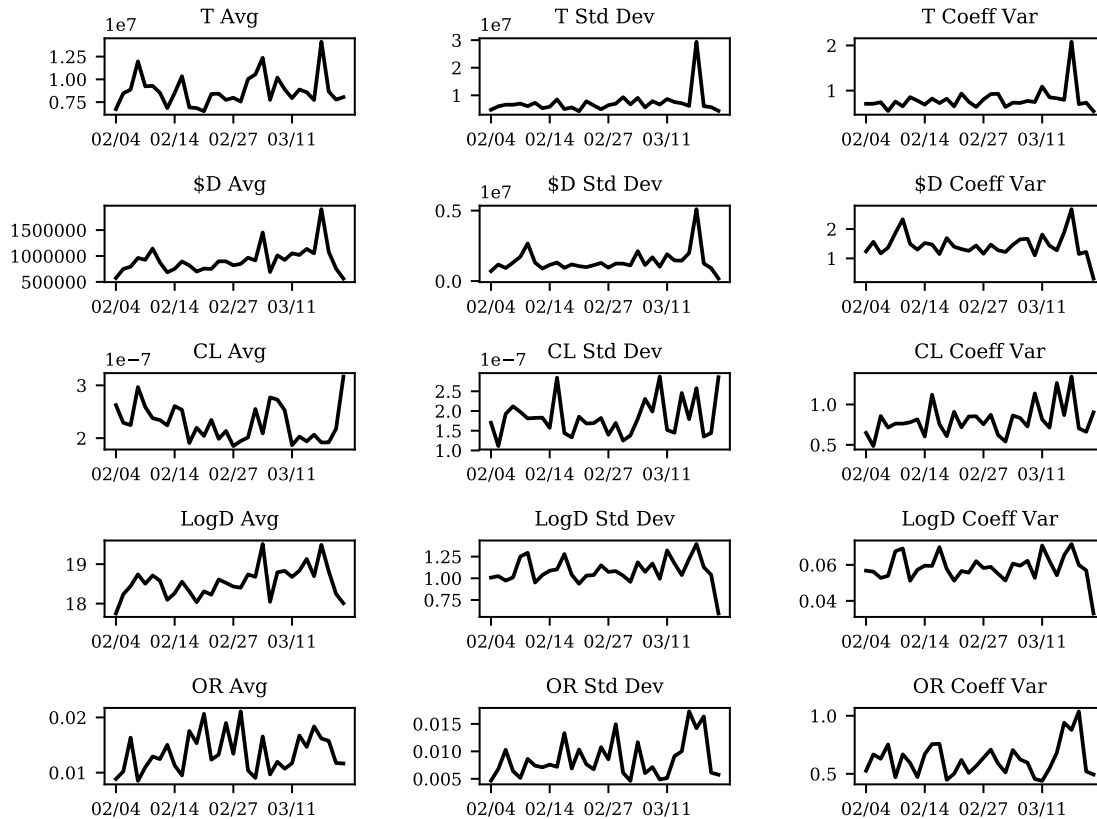


Figure 4-7: The average, standard deviation, and coefficient of variation of Citigroup’s volume and depth measures, aggregated across 2,485 five-minute bars.

other measures do. One might speculate that this dimension of liquidity may be more short-term as it experiences quite a bit of fluctuation across the time period.

Concerning trade count measures, as shown in Table 4.7, the average number of trades during each time bar is 850 with a minimum of 220 and a maximum of 1,610. The flow ratio, which is the ratio of turnover to average trade waiting time is on average \$41,400,000. Because this flow ratio is significantly greater than the average turnover, it indicates that the average wait time in between trades is less than a second. Together the large trade count and flow ratio suggests that volume traded mostly occurs over series of small trades rather than large, block trades.

However, there does appear to be some variation in this. Both of these measures are significantly skewed to the right, which illustrate that, like volume and depth, when there is significant variation in these measures, it is the result of a large increase

	N	FR
Mean	8.50e+02	4.14e+07
Median	6.88e+02	1.55e+07
Max	1.61e+04	1.17e+10
Min	2.20e+02	1.43e+06
Std Dev	6.60e+02	2.52e+08
Coeff Var	0.78	6.09
Skewness	7.79	40.12
Kurtosis	128.59	1834.1

Table 4.7: Summary statistics of Citigroup’s trade count measures; these encompass 2,485 five-minute bar observations for each measure.

of trades or high liquidity, rather than a decrease and low liquidity. The kurtosis for flow ratio is incredibly large, signaling that there is a great amount of dispersion in the tails for this measure. However, given this, the coefficient of variation for flow ratio at 6.09 seems rather low.

Figure 4-8 perhaps better illuminates the dispersion of these measures as it shows their daily averages, standard deviations, and coefficients of variation over the time period of analysis. While there are many peaks in trade count and an obvious sharp increases on in both trade count and flow ratio on March 15, 2019, relative to other points in time, the flow ratio’s reaction is more drastic than it might normally be perceived to have been. This is shown by the magnitude of the increases in flow ratio’s standard deviation and coefficient of variation, compared to the rest of the time period. Because the the flow ratio should be directly proportional to the trade count (because it is essentially divided by the inverse of trade count) and given that it reacted more aggressively than the trade count did, this suggests that there was even more variation in the turnover that day, relative to trade count.

The statistics for Citigroup’s price change measures are displayed in 4.8. These measures appear to have incredibly small values across all statistics which can be explained by the fact that all of these measures are divided by turnover or average trade wait time, which as shown above, are very large. Perhaps because the values themselves are so small and are more ratios than anything else, it is useful to look at the way these ratios change over the 32-day time period.

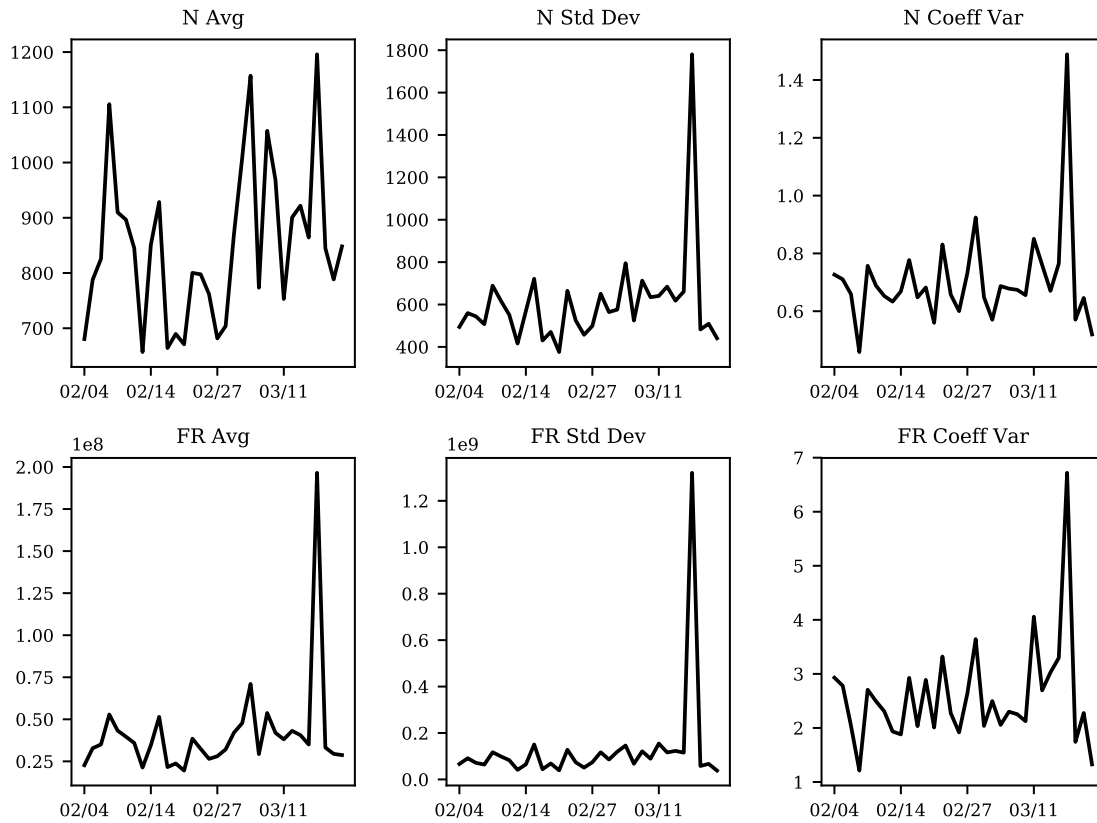


Figure 4-8: The average, standard deviation, and coefficient of variation of Citigroup's trade count measures, aggregated across 2,485 five-minute bars.

	M	LR1	Illiq
Mean	8.45e-04	5.58e-08	1.19e-10
Median	1.39e-05	4.78e-08	9.73e-11
Max	2.41e-01	2.05e-06	1.94e-09
Min	1.21e-06	2.32e-09	0.00e+00
Std Dev	8.15e-03	5.77e-08	1.03e-10
Coeff Var	9.65	1.03	0.87
Skewness	20.16	0.0	0.0
Kurtosis	493.09	0.0	0.0

Table 4.8: Summary statistics of Citigroup's liquidity measures involving price changes; these encompass 2,485 five-minute bar observations for each measure.

The coefficient of variation, skewness, and kurtosis values for the Martin Index are all quite significant. From this, the Martin Index is right-skewed, which together with the large tails noted by the kurtosis value, indicates that the Martin Index has quite a meaningful amount of dispersion to more positive values. This is perhaps somewhat surprising as for the most part, the other measures' skewness is in the direction of more liquidity. However, like spread measures, a low value correspond to high liquidity, so the Martin Index is dispersed more significant over low liquidity levels.

Alternatively, liquidity ratio one and ILLIQ have small coefficients of variation at 1.03 and 0.87 respectively and a skewness and kurtosis of zero. Perhaps this is more of a function of the very small range of values that liquidity ratio one and ILLIQ take because of their large denominators rather than a statement on their dispersion.

Figure 4-9 illustrates the daily average, standard deviation, and coefficient of variation of these measures for the entirety of the time period. The Martin index presents the most readable peaks as its daily average seems to vary quite a bit. However, its coefficient of variations indicates a more stable pattern in a sense as it remains in the same range with only slight variation and mostly centers around a coefficient of four. Meanwhile, the daily averages for ILLIQ and liquidity ratio one center around a mean value for the most part, they indicate unusual values in their measures on two specific days; this can be seen through the dramatic single value in the standard deviation and coefficient of variation.

Although the measures seem quite similar, Figure 4-9 suggests that they measure quite different dimensions of liquidity as their values and standard deviations significant spike at different time points. Furthermore, because the coefficient of variation of the Martin Index relatively centers around four, albeit with quite amount of variation, this suggests that the Martin Index might experience mean reversion over a certain time period or captures some measure that is relatively stationary across the time series.

Given the sense of how these measures behave over a 32-day time period, we can examine their aggregate 15-minute Z-scores, which were created by normalizing

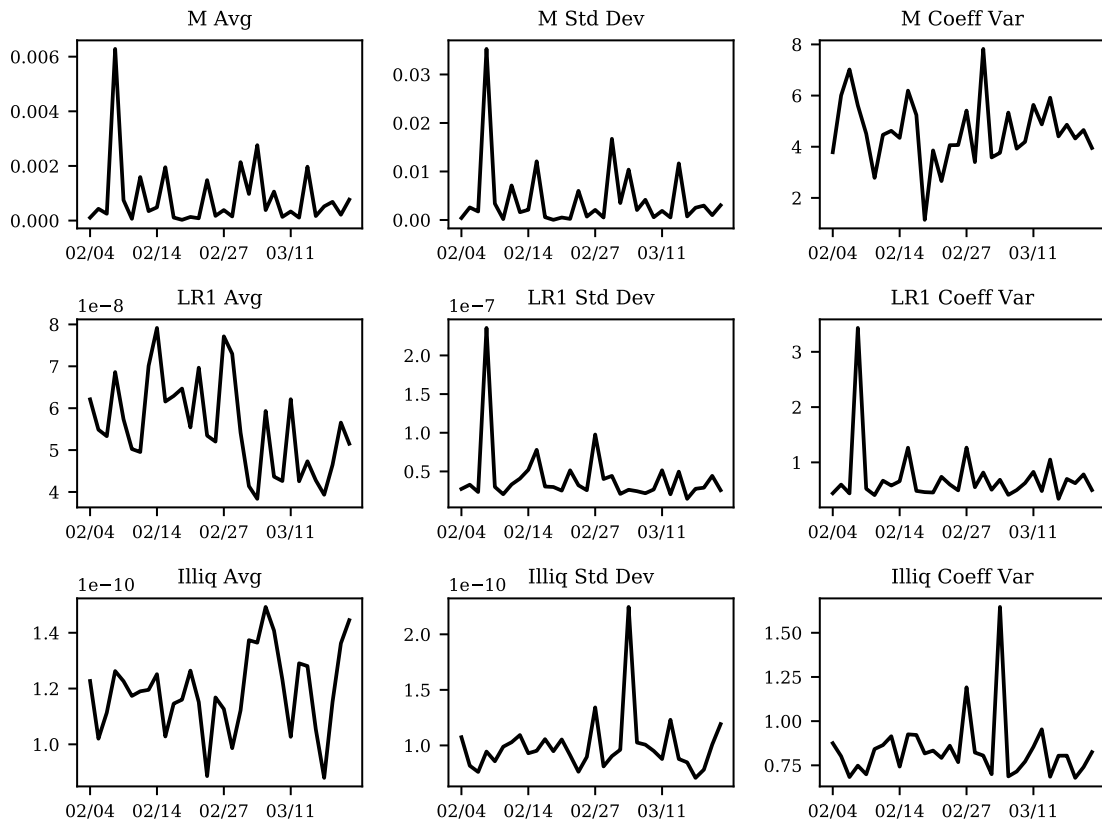


Figure 4-9: The average, standard deviation, and coefficient of variation of Citigroup's price-change measures, aggregated across 2,485 five-minute bars.

each measure by each day, and then averaging those Z-scores across 15-minute bars across all 32 days in the sample. Figure 4-10 displays the intraday 15-minute average Z-scores across the time period for each liquidity measure.

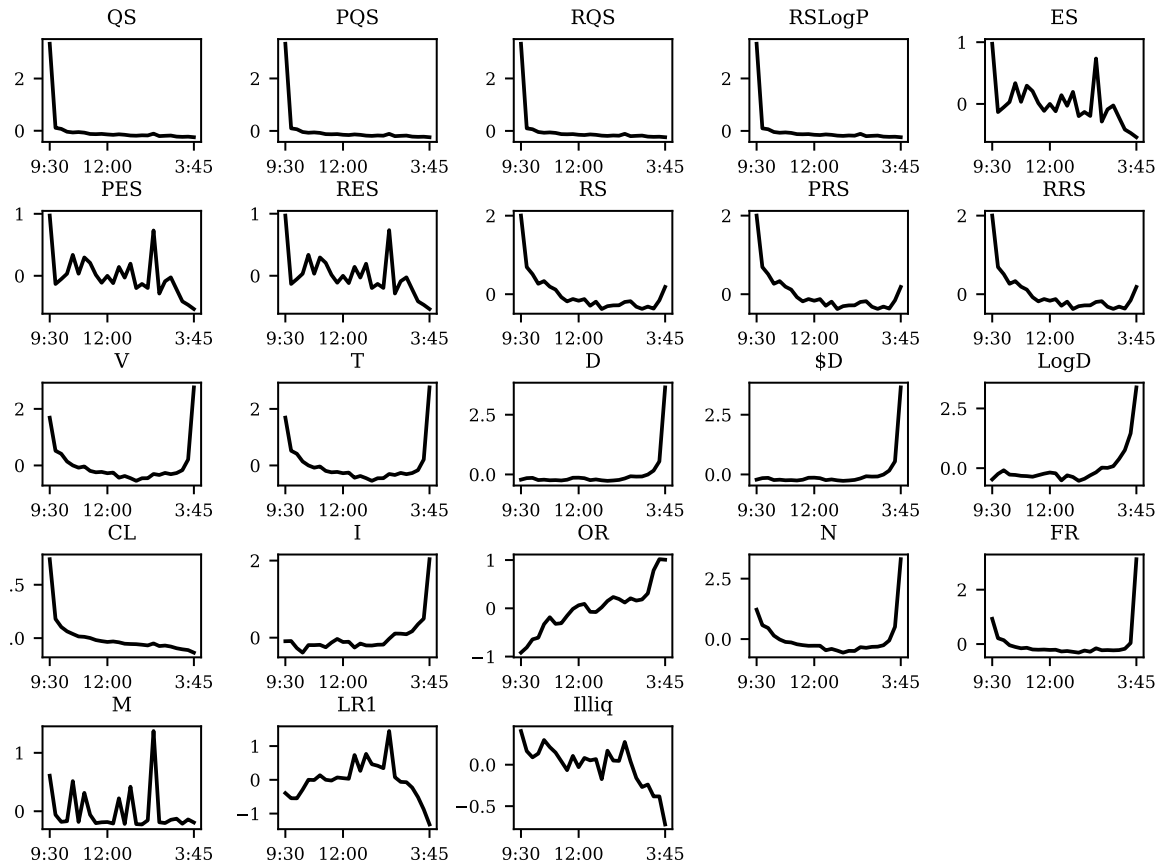


Figure 4-10: The Z-scores of Citigroup's liquidity measures averaged across 15-minute intervals.

On an intraday basis, the spread measures, shown in Figure 4-11, do not seem to exhibit the inverted U-shape one might expect. The quoted and realized spreads, and to a less extent the effective spread, show greater spreads (or rather lower liquidity) in the morning and then a decrease throughout the day. The realized spread does slightly increase at the end of its day, sort of forming a lopsided U-shape. This suggests that price movement from mid after five minutes increases in the morning and close to market close. Because the effective spread is quite variable and seems to jump from the mean of one 15-minute bar to the next, it might suggest that the values from a few particular days are overpowering the averages across the 32 days. This possible

explanation is supported by referring back to Figure 4-6, which shows there are a few days where the effective spread is quite large relative to others. Although the realized spread has the greatest coefficient of variation, the variation is more evenly distributed across the days, providing a plot that does seem to demonstrate an intraday liquidity pattern.

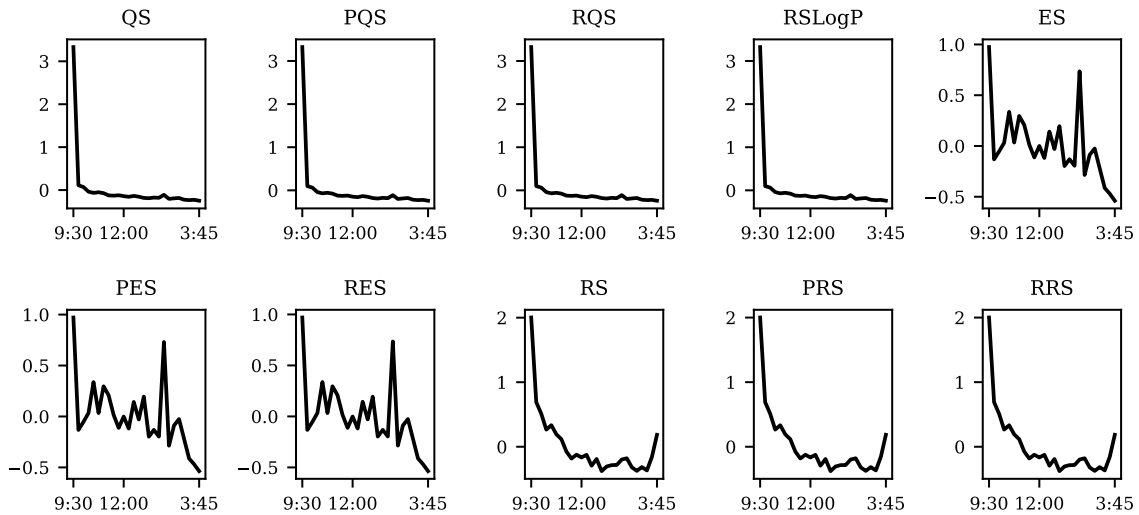


Figure 4-11: The Z-scores of Citigroup's spread measures averaged across 15-minute intervals.

As for the volume and depth measures, shown in Figure 4-12, the depth measures do mirror and support the intraday patterns exhibited by the spread measures. Depth, dollar depth, and log depth all are very low in the morning hours and then increase, reaching its most liquid point in the day near market close. The fact that quoted spreads and quoted depths mirror each other in intraday liquidity patterns perhaps indicates that these measures move in tandem (i.e. depth or spread are being adjusted together rather than one or the other in times of both higher and lower liquidity).

In fact the composite liquidity measures this to some degree, as it is the ratio of the proportional quoted spread to the dollar depth, as it measures the slope of the quotes. However, if the proportional quoted spread and dollar depth were adjusted in perfect unison where during time of low (high) liquidity, the quoted spread increases




Figure 4-12: The Z-scores of Citigroup's volume and depth measures averaged across 15-minute intervals.

(decreases) and the depth decreases (increases), than the composite liquidity would be constant throughout the day. So while the shape of composite liquidity's intraday pattern further affirms the notion that there is low liquidity in the morning and high liquidity near the market close in terms of quotes, it also suggests the quoted spread and dollar depth move in somewhat of a disjoint way. The derivative of the composite liquidity would provide more insight into this behavior.

Figure 4-13 displays the intraday liquidity patterns for trade count and the flow ratio, which is the ratio of turnover to trade waiting time. Both of these display exactly the same pattern, a U-shape that is skewed upward near the end of the day. This indicates these measures are most liquid in the mornings and near market close, but market close is still more liquid than the morning is. While the flow ratio, as shown in Table 4.7 and Figure 4-8, demonstrated greater variation on across the time period, trade count exhibits greater variation on an intraday basis as indicated by the larger range it spans (about -0.5 to just under 3.5) compared to that of the flow ratio (roughly -0.3 to 3.1).

While the trade count measures for Citigroup demonstrated fairly consistent in-




Figure 4-13: The Z-scores of Citigroup's trade count measures averaged across 15-minute intervals.

trading patterns, the same cannot be said for its price change measures. Figure 4-14 shows these intraday patterns, where the lack of smoothness suggests that the average Z-scores for a few specific days may be skewing the averages across days. However, if we were to extrapolate insight from this, we might tentatively say that ILLIQ decreases over the course of the day, which supports the notion of increasing liquidity as the trading day goes on. Liquidity ratio one appears to roughly be an inverted U-shape, which indicates that the average price change of a transaction is greatest during the middle of the day, slightly less in the morning, and the lowest near market close, all supporting the U-shape with the slightly larger tail near market close demonstrated by volume, turnover, trade count, and flow ratio.

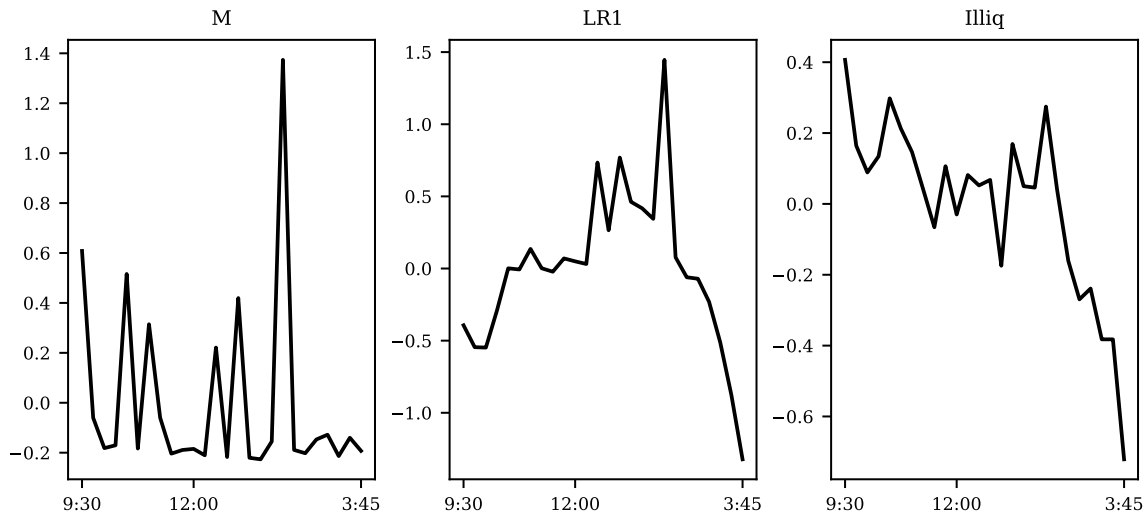


Figure 4-14: The Z-scores of Citigroup’s price change composite measures averaged across 15-minute intervals.

4.3 Liquidity Measures for the Financial Sector ETF

Tables 4.9 through 4.12 contain the statistics of the Financial Sector ETF’s liquidity measures on 2,496 aggregated five-minute time bars. Despite the difference in security type, as shown below, many of the observations that are consistent across Bank of America and Citigroup, particularly regarding the shape of intraday liquidity patterns for specific measures, are generally true here as well.

The statistics of the spread measures are in Table 4.9 and the average, standard deviation, and coefficient of variation of the quoted, effective, and realized spreads over the 32-day analysis period are displayed in Figure 4-15. The average quoted spread for a five-minute bar is \$0.01 and the average effective spread is \$0.004. The average effective spread being less than the minimum tick size indicates that trading at mid occurs the majority of the time, as the average effective spread is closer to \$0.00 than it is to the minimum tick size \$0.01. Meanwhile, the realized spread, which measures two times the absolute value of the difference between trade price and the

	QS	PQS	RQS	RSLogP	ES
Mean	1.00e-02	3.80e-04	3.80e-04	3.80e-04	4.37e-03
Median	1.00e-02	3.79e-04	3.79e-04	3.79e-04	4.27e-03
Max	1.10e-02	4.25e-04	4.25e-04	4.25e-04	1.54e-02
Min	1.00e-02	3.69e-04	3.69e-04	3.69e-04	0.00e+00
Std Dev	4.69e-05	5.35e-06	5.35e-06	5.35e-06	2.49e-03
Coeff Var	0.003	0.01	0.01	0.01	0.57
Skewness	10.76	0.97	0.98	0.97	0.58
Kurtosis	179.8	3.44	3.44	3.45	0.78
	PES	RES	RS	PRS	RRS
Mean	1.66e-04	1.66e-04	3.25e-02	1.24e-03	1.24e-03
Median	1.63e-04	1.62e-04	3.00e-02	1.12e-03	1.11e-03
Max	5.74e-04	5.74e-04	7.50e-01	2.79e-02	2.83e-02
Min	0.00e+00	0.00e+00	0.00e+00	0.00e+00	0.00e+00
Std Dev	9.46e-05	9.46e-05	3.39e-02	1.29e-03	1.29e-03
Coeff Var	0.57	0.57	1.04	1.04	1.05
Skewness	0.58	0.58	5.58	5.37	5.49
Kurtosis	0.76	0.76	85.95	79.37	83.02

Table 4.9: Summary statistics of the Financial Sector ETF's spread measures; these encompass 2,496 five-minute bar observations for each measure.

mid quote five-minute later, is \$0.03, which indicates about \$0.01-\$0.02 cents from the mid-quote to either the ask or the bid.

When examining the value of the average spread measures in Figure 4-15 over the course of the entire 32 days, there is some variation, but the ranges which these measures move in is quite small, as indicated by the small magnitude of their coefficients of variation.

These measures do not seem to move in unison with each other; from the coefficients of variation over the time period, it appears the quoted spread was more variable in the beginning of the time period than the end, suggesting some shift in overall liquidity. Meanwhile, the coefficient of variations for the effective spread vary quite a bit on each day but overall center around 0.55, or rather 0.57 as Table 4.9 indicates.

While the quoted spread may have demonstrated smaller liquidity which then stabilized in the beginning of the period, the effective spread does not appear to indicate any shift in liquidity over the time period. The realized spread, however,

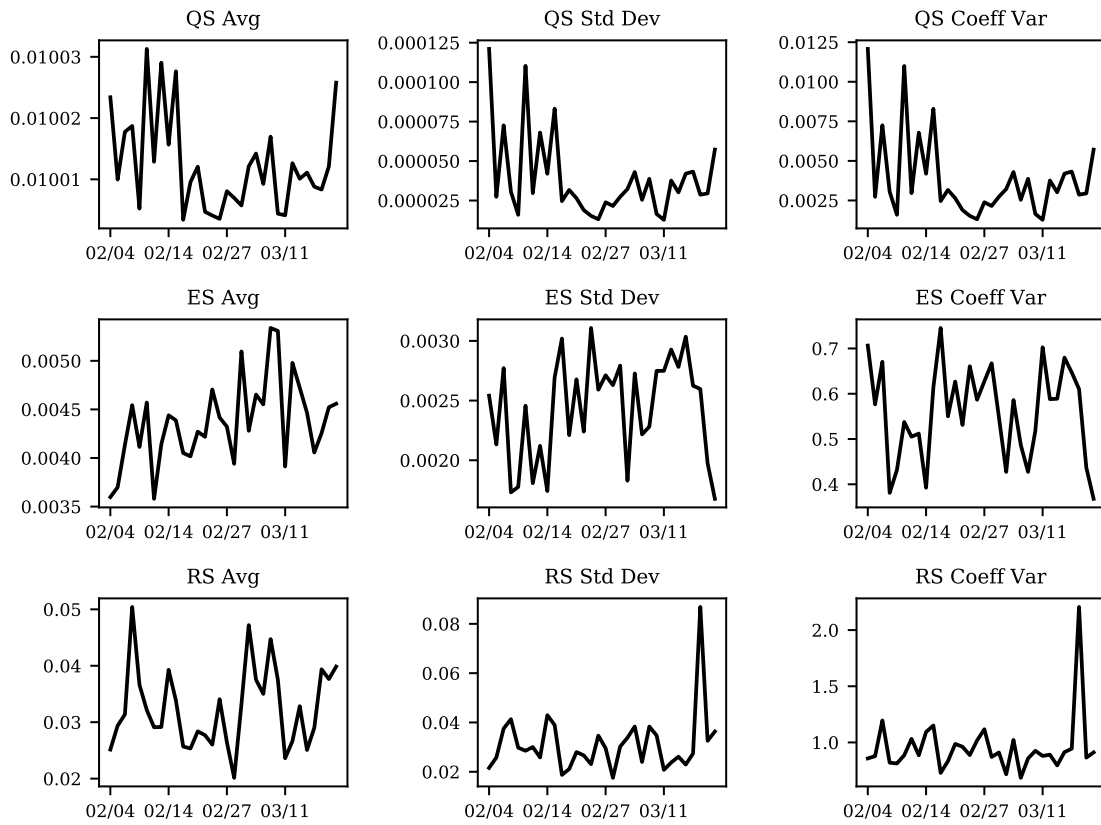


Figure 4-15: The average, standard deviation, and coefficient of variation of the Financial Sector ETF's spread measures, aggregated across 2,496 five-minute bars.

	V	T	D	\$D
Mean	5.08e+05	1.34e+07	3.55e+06	4.67e+07
Median	3.38e+05	8.84e+06	2.53e+06	3.33e+07
Max	7.69e+06	2.05e+08	4.89e+07	6.50e+08
Min	2.27e+04	5.99e+05	3.25e+05	4.32e+06
Std Dev	5.46e+05	1.43e+07	3.96e+06	5.22e+07
Coeff Var	1.07	1.07	1.12	1.12
Skewness	4.35	4.36	5.59	5.59
Kurtosis	32.03	32.46	42.33	42.46
	CL	LogD	I	OR
Mean	9.24e-09	2.82e+01	1.61e+07	3.10e+00
Median	6.87e-09	2.80e+01	9.77e+06	2.07e+00
Max	3.75e-07	3.40e+01	2.81e+08	3.45e+01
Min	3.48e-09	2.39e+01	4.60e+03	8.29e-02
Std Dev	1.65e-08	1.30e+00	2.26e+07	2.95e+00
Coeff Var	1.78	0.05	1.4	0.95
Skewness	0.0	0.83	5.07	2.84
Kurtosis	0.0	1.83	40.39	13.28

Table 4.10: Summary statistics of the Financial Sector ETF’s volume and depth measures; these encompass 2,496 five-minute bar observations for each measure.

centers around average values (both in terms of the measure’s values themselves, its standard deviations, and its coefficients of variation), until the peak on March 18, 2019, indicating a shock of low liquidity. While this isn’t immediately obvious from its daily average, it is indicated by its standard deviations and coefficients of variation.

The effect of the stability of these spreads is also demonstrated by their skew and kurtosis values. The spread measures are all skewed to the right, however some to not such a large degree. For instance, the proportional and relative quoted spreads, the relative spread of log prices, and all of the effective spread measures all demonstrate a skew of less than one.

The skewness for the quoted spread and the realized spread measures are much larger in the range from 5.37 for the proportional realized spread to 10.76 for the quoted spread. The effective spread measures have kurtosis values around 0.76-0.78 and as such are thin-tailed. The remaining spread measures have large tails. These results can mostly be explained by significant variation in a few days, compared to relative stability the remainder of the time period.

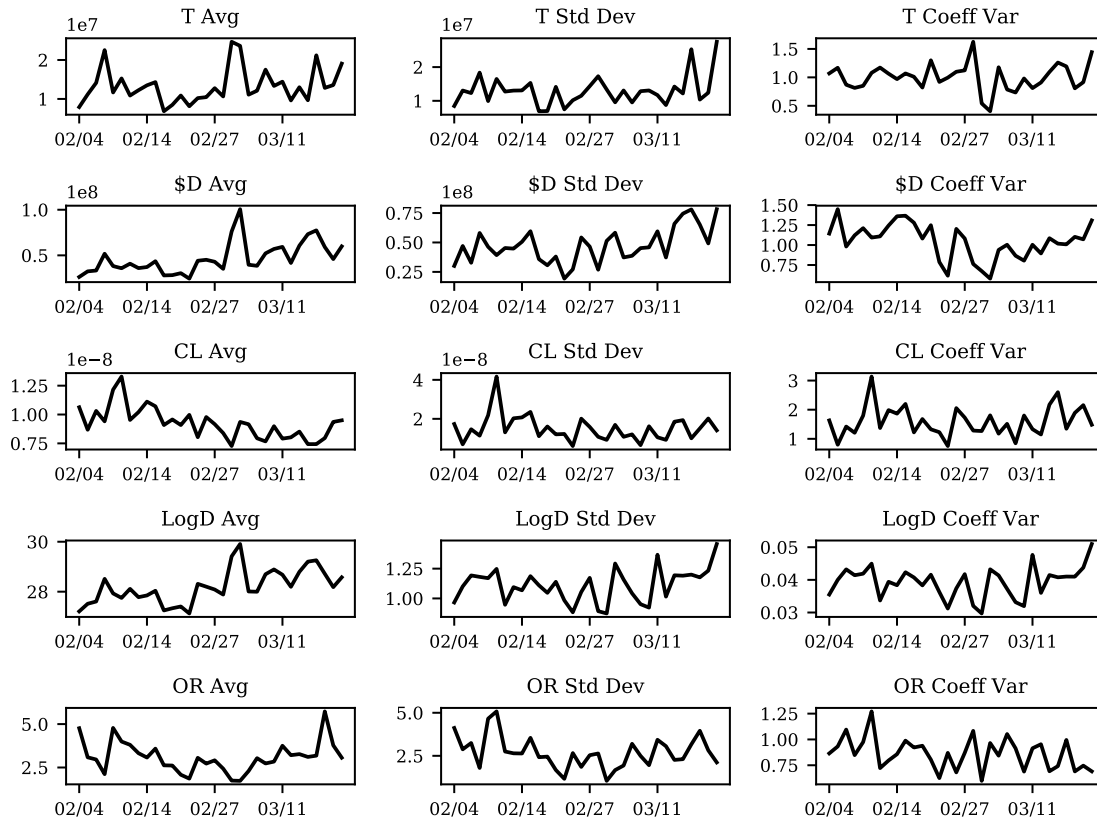


Figure 4-16: The average, standard deviation, and coefficient of variation of the Financial Sector ETF's volume and depth measures, aggregated across 2,496 five-minute bars.

For volume and depth measures, as indicated in Table 4.10, during the Financial Sector ETF's average five-minute bar, 508,000 shares, or rather \$13,400,000 in turnover is traded. Composite liquidity values are very close to zero because the dollar depth in the denominator of the measure is much greater than the proportional quoted spread. The average imbalance dollar amount between quoted bid and ask sizes is \$16,100,000 and 310% of the its turnover. This is an astoundingly high percent and indicates that there are significant differences between the bid and ask quote sizes compared to what is being traded in the market. These measures are right skewed and are fat tailed with the exception of log depth and composite liquidity which have thin tails.

Concerning trade count measures, as shown in Table 4.11, the average number of trades during each time bar is 805 with a minimum of 140 and a maximum of 8,370.

	N	FR
Mean	8.05e+02	6.27e+07
Median	6.30e+02	1.89e+07
Max	8.37e+03	5.71e+09
Min	1.40e+02	2.84e+05
Std Dev	6.62e+02	2.21e+08
Coeff Var	0.82	3.52
Skewness	4.42	13.93
Kurtosis	29.89	277.42

Table 4.11: Summary statistics of the Financial Sector ETF's trade count measures; these encompass 2,496 five-minute bar observations for each measure.

The flow ratio, which is the ratio of turnover to average trade waiting time is on average \$62,700,000 per a second. Because this flow ratio is significantly greater than the average turnover, it indicates that the average wait time in between trades is less than a second. Together the large trade count and flow ratio suggests that volume traded mostly occurs over series of small trades rather than in block trades.

However, there does appear to be some variation in this. Both of these measures are skewed to the right, which illustrate that, like volume and depth, when there is significant variation in these measures, it is the result of a large increase of trades or high liquidity, rather than a decrease and low liquidity. The kurtosis for flow ratio is quite large, signaling that there is a great amount of dispersion in the tails for this measure. However, given this, the coefficient of variation for flow ratio, 3.52, seems rather low.

Figure 4-17 perhaps better illuminates the dispersion of these measures as it shows their daily averages, standard deviations, and coefficients of variation over the time period of analysis. While there are many peaks in trade count and an obvious sharp increases in both trade count and flow ratio on March 15, 2019 and March 20, 2019, relative to other points in time, the flow ratio's reaction seems quite drastic.

This is shown by the magnitude of the increases in flow ratio's standard deviation and coefficient of variation, compared to the rest of the time period. Because the the flow ratio should be directly proportional to the trade count (because it is essentially divided by the inverse of trade count) and given that it reacted more aggressively

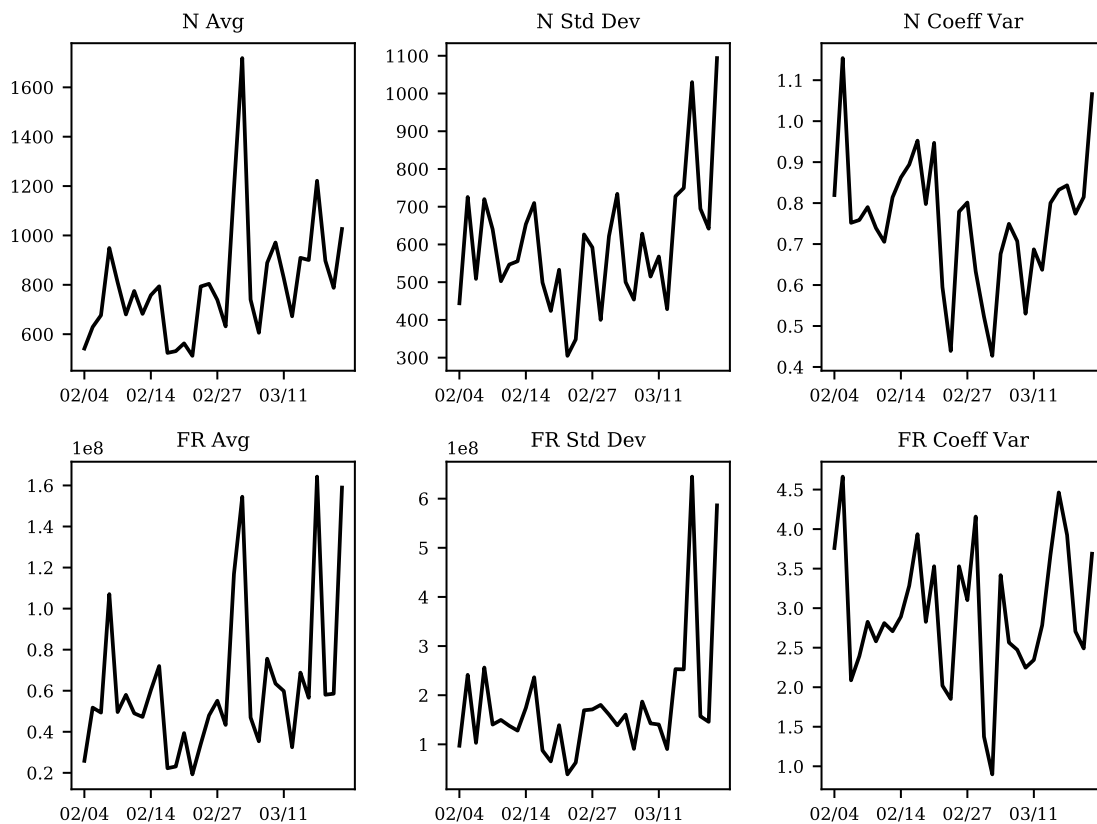


Figure 4-17: The average, standard deviation, and coefficient of variation of the Financial Sector ETF's trade count measures, aggregated across 2,496 five-minute bars.

	M	LR1	Illiq
Mean	9.73e-05	8.71e-08	6.91e-11
Median	6.91e-06	7.19e-08	4.97e-11
Max	2.74e-02	7.56e-07	2.86e-09
Min	5.86e-07	4.88e-09	0.00e+00
Std Dev	8.22e-04	6.37e-08	8.98e-11
Coeff Var	8.44	0.73	1.3
Skewness	19.44	0.0	0.0
Kurtosis	534.89	0.0	0.0

Table 4.12: Summary statistics of the Financial Sector ETF’s liquidity measures involving price changes; these encompass 2,496 five-minute bar observations for each measure.

than the trade count did, this suggests that there was even more variation in the turnover that day, relative to trade count.

The statistics for price change measures are displayed in 4.12. These measures appear to have incredibly small values across all statistics which can be explained by the fact that all of these measures are divided by turnover or average trade wait time, which both, as shown above, are very large. Perhaps because the values themselves are so small and are more ratios than anything else, it is useful to look at the way these ratios change over the 32-day time period.

The coefficient of variation, skewness, and kurtosis values for the Martin Index are large, indicating that the Martin Index has quite a meaningful amount of dispersion to more positive values. Like spread measures, a low value correspond to high liquidity, so the Martin Index is dispersed more significant over low liquidity levels. Liquidity ratio one and ILLIQ meanwhile have small coefficients of variation at 0.73 and 1.3 respectively and a skewness and kurtosis of zero. Perhaps this is more of a function of the very small range of values that liquidity ratio one and ILLIQ take because of their large denominators rather than a statement on their dispersion.

Figure 4-18 illustrates the daily average, standard deviation, and coefficient of variation of these measures for the entirety of the time period. The Martin index presents the most readable peaks as its daily average seems to vary quite a bit. However, its coefficient of variations indicates a more stable pattern in a sense as it remains in the same range with only slight variation and mostly centers around a

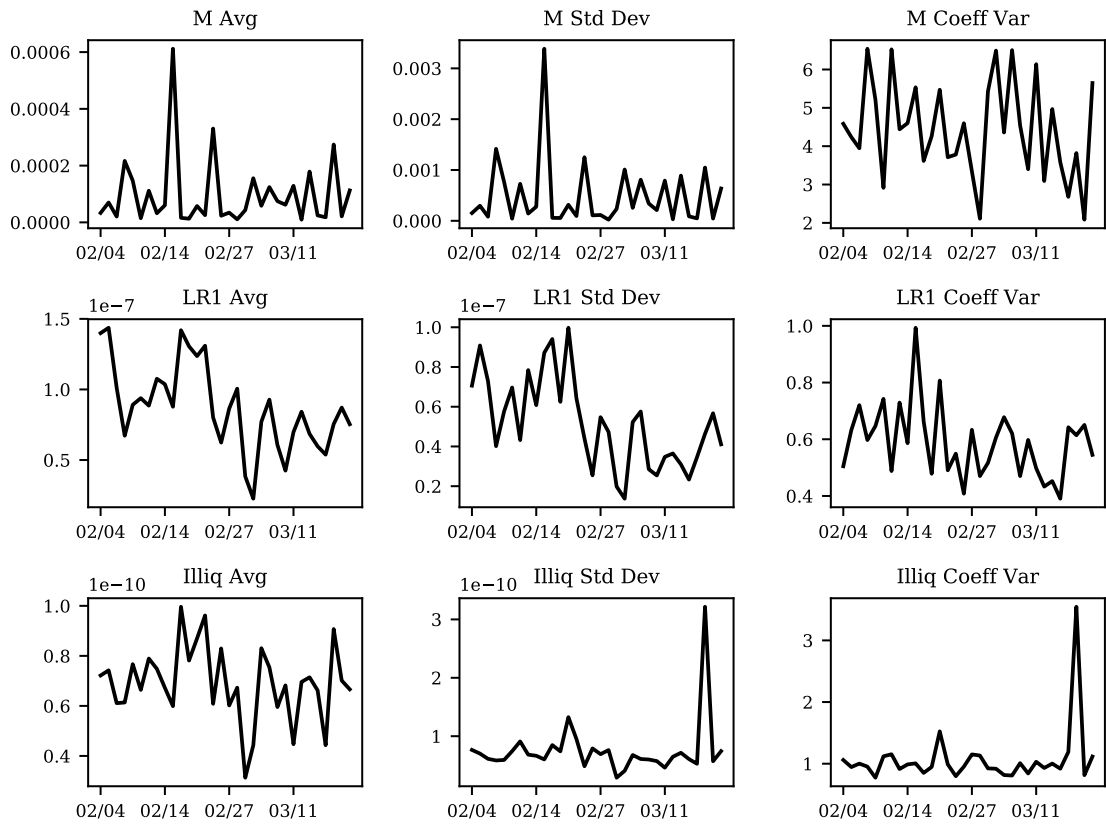


Figure 4-18: The average, standard deviation, and coefficient of variation of the Financial Sector ETF's price-change measures, aggregated across 2,496 five-minute bars.

coefficient of four.

While the daily averages for liquidity ratio one center around a mean value for the most part, just as with the Martin Index, its coefficients of variation center around 0.6, with perhaps a slight downward trend in these coefficients over the course of the time period. Alternatively, ILLIQ behaves in the opposite way; its standard deviations and coefficients of variation are relatively stable throughout the time period until March 18, 2019, when its standard deviation and consequently coefficient of variation dramatically increase. This difference in behavior perhaps indicates that ILLIQ captures a dimension of liquidity that the Martin Index and liquidity ratio one do not.

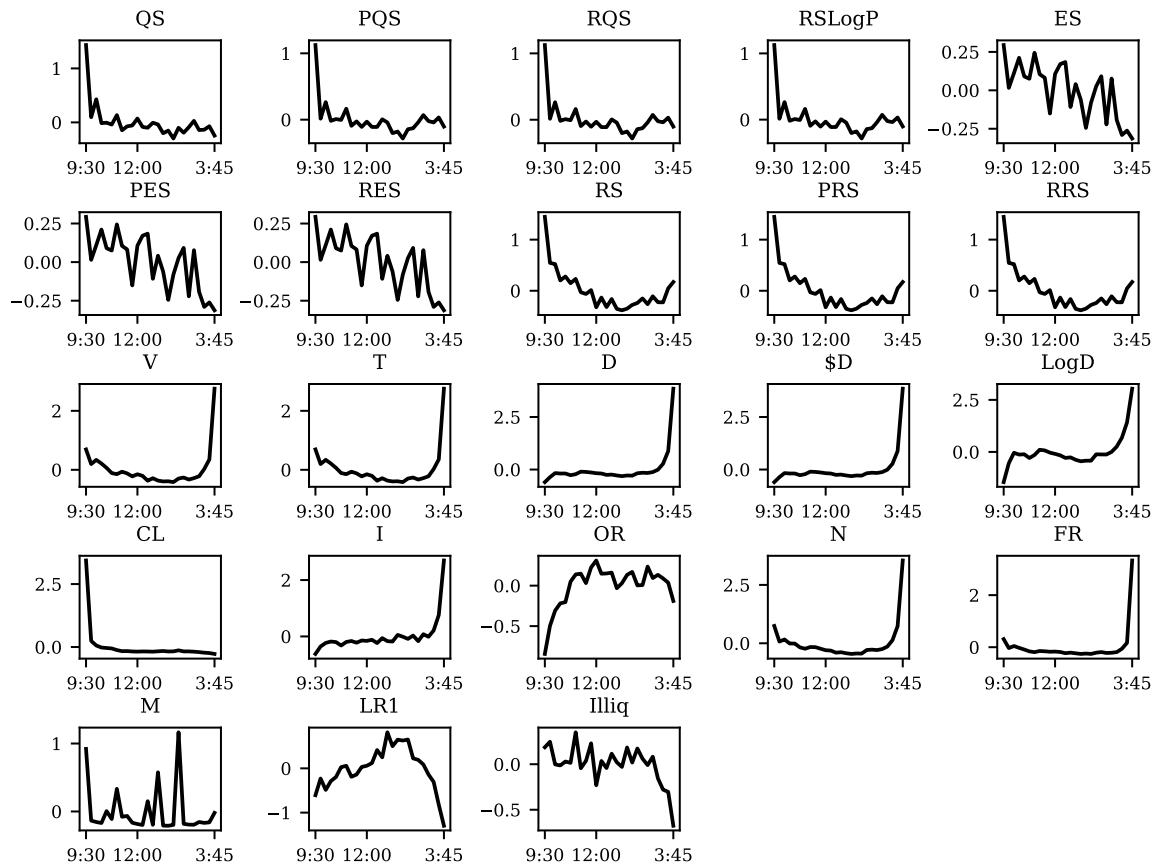


Figure 4-19: The Z-scores of the Financial Sector ETF's liquidity measures averaged across 15-minute intervals.

Given the sense of how these measures behave over a 32-day time period, we can examine their aggregate 15-minute Z-scores. Figure 4-19 displays the intraday

15-minute average Z-scores across the time period for each liquidity measure.

On an intraday basis, the spread measures, shown in Figure 4-20, follow a similar pattern as those seen by Bank of America, and even more closely, Citigroup. The quoted and realized spreads, and to a less extent the effective spread, show greater spreads (or rather lower liquidity) in the morning and then a decrease throughout the day. The realized spread does slightly increase at the end of its day, sort of forming a lopsided U-shape. This suggests that price movement from mid after five minutes increases in the morning and near market close.

Because the effective spread is quite variable and seems to jump from the mean of one 15-minute bar to the next, it might suggest that the values from a few particular days are overpowering the averages across the 32 days. This possible explanation is supported by referring back to Figure 4-15, which shows there are a few days where the effective spread is quite large relative to others.

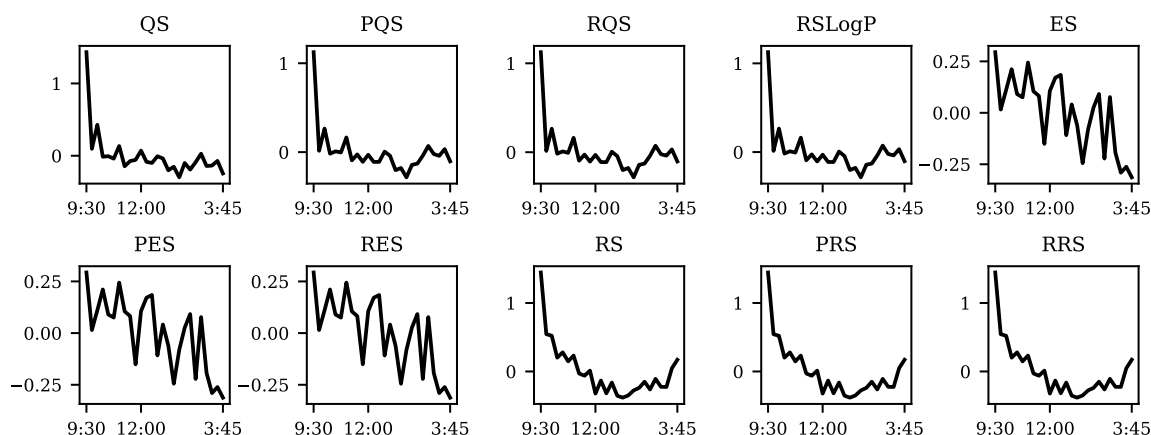


Figure 4-20: The Z-scores of the Financial Sector ETF's spread measures averaged across 15-minute intervals.

In general, these averages across time periods appear to be quite noisy, although the general pattern is somewhat clear after having analyzed the results for Bank of America and Citigroup. Although the quoted spread has a smaller coefficient of variation, it spans a larger range on an intraday basis compared to the effective spread. However, it experiences more intraday variation as indicated. Oddly this is opposite of the conclusions from the kurtosis and skew. This suggests, the effective spread has

relatively little variation over time but not on an intraday basis.

As for the volume and depth measures, shown in Figure 4-21, the depth measures do mirror and support the intraday patterns exhibited by the spread measures. Depth, dollar depth, and log depth are all low in the morning hours and then increase, reaching its most liquid point in the day near market close. The shape of the log depth intraday pattern is somewhat interesting and suggests that depth increases more quickly in the morning to reach a stable depth level, and then increases less quickly near market close to reach its peak amount.




Figure 4-21: The Z-scores of the Financial Sector ETF's volume and depth measures averaged across 15-minute intervals.

The shape of the imbalance of bid and ask depth indicates that there are not large discrepancies in the quoted sizes until later in the day and that this does appear to be a systematic pattern. Meanwhile, the order ratio, which describes the absolute value of the difference in bid and ask quote sizes divided by turnover, indicates that liquidity is highest in the morning and then decreases throughout the day. The order ratio begins to rise in the morning when the turnover decreases, as the depth spread appears stable (as indicated by the plot of the imbalance between buy and sell bid ask sizes). It stabilizes in the middle of the day when the depth spread still remains the same and turnover has stabilized at its lower middle of the day value.

Finally, the imbalance spread increases at the end of the day as does turnover. The imbalance between quoted sizes begins to increase earlier than the turnover does,

which is why the order ratio remains at a high value even near the end of the day. It is not until the turnover dramatically increases in the last 30 minutes of the day that the order ratio decreases accordingly.




Figure 4-22: The Z-scores of the Financial Sector ETF's trade count measures averaged across 15-minute intervals.

Figure 4-22 displays the intraday liquidity patterns for trade count and the flow ratio, which is the ratio of turnover to trade waiting time. Both of these display exactly the same pattern, a U-shape that is skewed upward near the end of the day and begins at a relatively low point in the beginning of the day. This indicates these measures are most liquid in the mornings and near market close, but market close is still more liquid than the morning is. While the flow ratio, as shown in Table 4.11 and Figure 4-17, demonstrated greater variation on across the time period, trade count exhibits slightly greater variation on an intraday basis as indicated by the larger range it spans compared to that of the flow ratio.

While the price change measures seemed to experience much variation over the time period without much of a clear trend (Figure 4-18), liquidity ratio one and ILLIQ seem to exhibit an intraday pattern of liquidity that is easier to decipher. For liquidity ratio one, the average change in price of a trade, a smaller price changes indicate greater liquidity. Thus, while there is clearly inter-day noise present in the

measure, liquidity ratio one's general intraday inverted U-shape supports the U-shape pattern of liquidity that previous measures have shown. ILLIQ presents a less clear picture in the morning, clearly showing a fair amount of noise with the variation, but does seem to decidedly decrease, indicating greater liquidity, at the end of the day. The Martin Index, on the other hand, either does not exhibit intraday liquidity patterns, or the outliers from a few days' distort the shape, making it challenging to extract trends from.

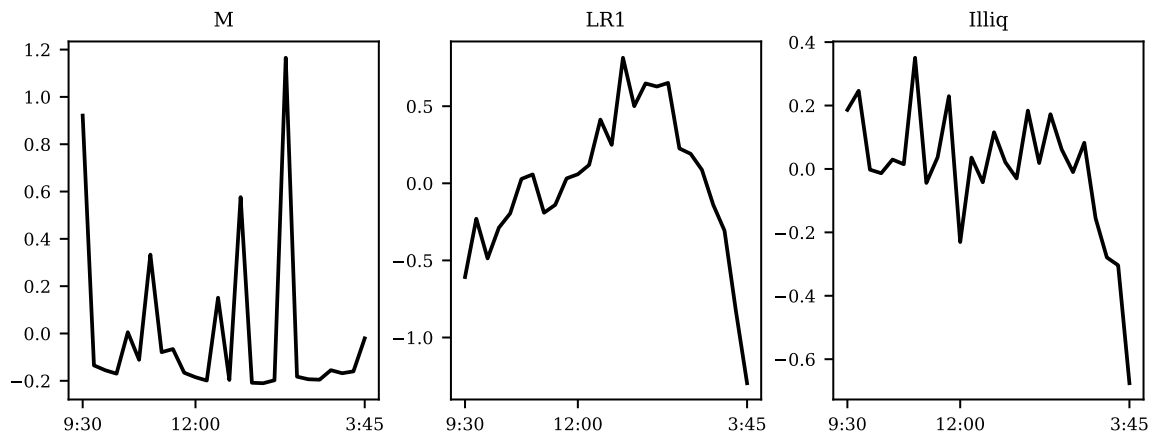


Figure 4-23: The Z-scores of the Financial Sector ETF's price change composite measures averaged across 15-minute intervals.

Chapter 5

Principal Components Analysis of Liquidity Measures

To examine common liquidity factors across liquidity measures for a given security, we perform a principal component analysis. While doing this analysis for several securities, we found that subsets of liquidity measures were highly correlated. This first became apparent in Chapter 4 when across all securities analyzed, the summary statistics for specific sets of variables were virtually indistinguishable. These measures were so highly correlated in fact that the variable loadings of these subsets in the top principal components that resulted from the analysis were effectively equivalent. Because of their equivalence, these subsets of measures are redundant and performing principal component analysis with all of these measures introduces the risk that measures are double or even triple counted in some cases. For this reason, we exclude the following measures from our principal component analysis: quoted spread, relative quoted spread, relative spread of log prices, effective spread, relative effective spread, realized spread, relative realized spread, volume, and depth. When choosing which measures to retain, we selected measures that did not rely on absolute price or volume to maximize transferability of measurements between securities. If after applying this criteria multiple measures per a subset still existed, we kept the measure that is most prevalent in literature.

An overview of the mechanics of principal component analysis is in section 5.1. Of primary interest is the loadings matrix which defines the coordinate axes of the principal components variables. It is common practice in factor analysis and principal components analysis to consider applying the “varimax” rotation to the loadings matrix, with the objective of defining interpretable factor variables that depend mostly on a few of the original variables. Section 5.2 reviews the varimax rotation. Sections 5.3 through 5.5 contain the analysis for each individual security.

5.1 Principal Components Analysis Theory

Principal components analysis (PCA) is a technique used to reduce the complexity of a multivariate data set. It accomplishes this by transforming n observations of p variables (that may or may not be correlated) into p linear combinations of the original variables. These p linear combinations are called the principal components. The components are uncorrelated and each explains a portion of the variance of the original data set. The principal components are ordered in such a way that the first component has the maximal variance; the second has maximal variance subject to being uncorrelated with the first; and so forth. Reducing the data set to principal components has two benefits: (1) only a fraction of the components will be deemed significant which reduces the dimensionality of the data set to be analyzed and (2) the components are uncorrelated and therefore allow for the results to be more easily interpreted.

Mathematically, PCA is an orthogonal linear transformation of the data to a new coordinate system. Consider the matrix X which consists of n rows (observations) and p columns (variables). If the scale of these variables varies greatly, it is standard to scale this input matrix so that each variable has a mean of zero and variance of one. After applying PCA, p new variables, known as principal components or factors, are constructed as weighted averages of the original p variables. Each of the n rows now has p new values corresponding to the principal components; these row values are known scores. Let Y be the score matrix. More formally, the equation for PCA

is:

$$Y = XW$$

where W is a matrix of coefficients that is determined by PCA. This equation may also be thought of as a set of p linear equations which weight the original variables to form each factor. As a linear equation, each score y_{ij} is:

$$y_{ij} = x_{i1}w_{j1} + x_{i2}w_{j2} + \dots + x_{ip}w_{jp}$$

W is an orthogonal matrix where its columns are orthogonal unit vectors. The first principal component variable maximizes variances subject to the first column of W being a unit vector. The second principal component maximizes variances subject to the constraint that the second column of W is a unit vector and that the covariance between the first and second columns of W is zero. This process continues for each of the p columns of W . We can calculate W using singular value decomposition (SVD). The covariance matrix, S , is defined as:

$$s_{ij} = \frac{\sum_{k=1}^n (x_{ki} - \bar{x}_i)(x_{kj} - \bar{x}_j)}{n}$$

Note that if the original variables are centered with a mean equal to zero then

$$S = \frac{1}{n}X^T X$$

S is a $p \times p$ symmetric matrix. Its diagonal values are the variances of each original variable, and the off-diagonals are the covariances between variables. Because S is symmetric, its SVD can be written as:

$$S = W' \Lambda W$$

where W is a matrix of the eigenvectors of S and Λ is a diagonal matrix whose diagonal elements are the eigenvalues corresponding to the eigenvectors of S . Therefore W can

be found by computing the normalized eigenvectors of S and arranging them such that the first column of W is the eigenvector with the largest eigenvalue, the second column is the eigenvector with the second largest eigenvalue, and so forth. The correlation between the i^{th} factor and the j^{th} original variable comes from:

$$r_{ij} = \frac{w_{ij}\sqrt{\lambda_i}}{\sqrt{s_{jj}}}$$

where w_{ij} is an element of w , λ_i is a diagonal element of λ , and s_{jj} is a diagonal element of S . These correlations are the component loadings.

There are a number of methods (both qualitative and quantitative) that are commonly used to determine the number of significant principal components. The scree test is a popular visual approach; it suggests retaining all factors to the left of the "elbow" on a scree plot. Scree plots display the eigenvalues from largest to smallest against the principal components. One of the most widely used analytical tests is applying the Kaiser criterion. The Kaiser criterion can be applied to the principal components derived from the correlation matrix of the original data. The criterion states that only the principal components with eigenvalues greater than one should be kept. The rationale behind this is that a factor is significant as long as it extracts at least as much information as the equivalent of one original variable. Because the data is centered and scaled in the correlation matrix, each original variable has a variance of one. A principal component's variance is equal to its eigenvalue. For our analysis, we use the Kaiser criterion to determine which components to retain. We use this for its widespread use in existing literature and its ease of interpretation.

5.2 Varimax Rotation of Principal Component Loadings

Principal component loadings, especially when a large number of variables are involved, can sometimes be difficult to interpret. A varimax rotation is a technique that can be applied to the significant principal components to further simplify the

results. This is accomplished through a change of coordinates such that the sum of the variances of the squared loadings is maximized. The goal of varimax is to produce a result where each factor has a small number of large loadings and a large number of zero or very small loadings. It has been found that maximizing the sum of variance of squared loadings tends to produce such features. As such, we apply a varimax rotation to the results from our PCA to aid the interpretation and understanding of the PCA.

5.3 Principal Components Analysis of Bank of America's Liquidity Measures

The principal component analysis for Bank of America is performed on 2,483 observations of 13 liquidity measures, or rather a data matrix X of dimension $2,483 \times 13$. The amounts of variance explained by the first five principal components are shown in Table 5.1. At 31%, the first component accounts for more than twice as much variance as any other component. The second and third components account for similar amounts of variances as do the fourth and fifth. Similar to a step function, this behavior can be seen in the scree plot of the eigenvalues of the principal components in Figure 5-1.

While the Kaiser criterion classifies five of the principal components as significant, the results from the scree plot are naturally more ambiguous than the analytical Kaiser criterion. One might interpret the elbow of the scree plot to be at component five or component seven in which case four or six components respectively would be significant. Given the results of applying the Kaiser criterion and the scree plot, designating the first five principal components as significant is reasonable.

The loadings produced from the PCA can be found in Figure 5-2 and Table 5.2. The first component largely captures the weighted average of the turnover, dollar depth, logarithmic depth, trade count, and flow ratio, which is the ratio of turnover to trade waiting time. For the most part, these are one-dimensional liquidity measures

	Standard deviation	Variance	Proportion of Variance	Cumulative Proportion
PC1	2.008338	4.033422	31.03%	31.03%
PC2	1.330799	1.771026	13.62%	44.65%
PC3	1.230438	1.513977	11.65%	56.30%
PC4	1.137215	1.293259	9.95%	66.24%
PC5	1.010667	1.021447	7.86%	74.10%

Table 5.1: The principal components of Bank of America using PCA on 13 liquidity measures and the Kaiser criterion to determine significance.




Figure 5-1: The scree plot of Bank of America's principal components resulting from PCA on 13 liquidity measures.

that concern the quantity of shares traded or quoted.

The second component accounts for 13.6% of the variance and is effectively an average of the the proportional effective spread and the Martin Index. The proportional effective spread measures the spread from the mid price at which each trade was executed as a proportion of its mid price; mainly it aims to quantify realized




Figure 5-2: The principal component loadings of Bank of America resulting from PCA on 13 liquidity measures.

	PC1	PC2	PC3	PC4	PC5
PQS	-0.011	-0.055	-0.164	-0.408	0.602
PES	-0.018	-0.674	-0.168	0.115	0.003
PRS	-0.021	0.053	-0.109	-0.002	-0.631
T	0.454	0.084	-0.205	0.039	0.011
\$D	0.448	-0.078	0.203	0.146	0.034
LogD	0.401	-0.121	0.286	-0.128	-0.014
CL	-0.053	-0.277	0.600	-0.029	0.244
OR	0.011	0.025	-0.540	-0.280	0.165
N	0.472	0.055	-0.127	0.040	0.035
FR	0.425	0.047	-0.104	0.170	0.066
M	-0.004	-0.629	-0.267	0.164	-0.103
LR1	0.089	-0.096	0.061	-0.618	-0.301
Illiq	0.133	-0.157	0.122	-0.513	-0.212

Table 5.2: The principal component loadings of each of Bank of America’s 13 liquidity measures across components.

transaction costs for an investor. The Martin Index measures the squared absolute price between consecutive trades and divides by the turnover.

Overall, this second factor describes the variance of price changes between trades and the spread from mid at which a trade occurs. Perhaps it is not surprising that these two measures are weighted equally as it suggests that when the spread from

the mid-price at which trades occur is wider (tighter), the price changes between the current trades and previous trades is greater (less). In a way, the value of the Martin Index is in part due to the behavior of the proportional effective spread.

The third component shows order ratio contrasted against composite liquidity. The order ratio is the absolute value of the difference between the sizes of the best bid and best ask quotes over the turnover. Note that because this measure uses the absolute value of the difference between the quoted bid and ask sizes, it is not directional. Composite liquidity is calculated as the proportional quoted spread over the average size of the best bid and ask prices.

Given these definitions, the third component might be thought of as the difference between how the bid and ask quote sizes are adjusted compared to how their spread is adjusted. The negative correlation of these measures initially suggests that as the imbalance between the sizes of the best bid and ask decreases (increases), the proportional quoted spread as a function of average quoted size increases (decreases). However, this dynamic is unexpected as one expects the depth imbalance to increase as the spread does.

As a possible explanation for this behavior, perhaps the negative correlation suggests that market makers maintain a balance between quoted sizes and spreads, consistently mitigating one with the other depending on the circumstance. Alternatively, it is also feasible that the negative correlation between order ratio and composite liquidity is the result of market makers decreasing their overall quoted sizes in times of illiquidity. This in turn would cause the absolute difference between the bid and ask sizes to decrease, which causes the order ratio to indicate higher liquidity in lower liquidity circumstances.

The fourth principal component explains 9.95% of the overall variance and can be described generally as a weighted average of liquidity ratio 1, the average price change of a transaction, and *ILLIQ*, the return over a five minute bar divided by the turnover during that bar. Through these two variables, the fourth component encompasses many of the dimensions of price change and quantity, including price changes of consecutive trades, trade count, price return over a longer term (the five-

minute bar), and turnover.

The fifth component demonstrates the contrast between the proportional quoted spread and the proportional realized spread, where the variables' loadings are of almost equal magnitude but negatively correlated. This component signifies the difference between liquidity risk priced in the market (as represented by the proportional quoted spread) and liquidity risk that is realized after five minutes. Here, the proportional realized spread and the quoted spread are directional. However, during the data cleaning process, all negative bid-ask spreads were removed, so the proportional realized spread is the only measure of the two that is capable of taking on negative values (when the mid price decrease five minutes after the trade).

While the negative correlation between these two measures is counterintuitive, perhaps it is an indication of market makers using the previous bar's price changes to adjust their quotes during the current bar. Their adjustment in the current bar could then in turn cause the realized spread's behavior to change, creating almost a self-perpetuating cycle of sorts. The analysis of component five's autocorrelation, discussed in detail in Chapter 6, supports this interpretation of the measure to a certain extent.

	PC1	PC2	PC3	PC4	PC5
PQS					0.742
PES		-0.7			
PRS					-0.526
LogD	0.369		0.28		
\$D	0.462				
N	0.476				
T	0.458				
CL			0.693		
M		-0.706			
LR1				-0.7	
FR	0.452				
OR			-0.485		0.382
Illiq				-0.598	

Table 5.3: A varimax rotation of Bank of America's first five principal component loadings resulting from principal components analysis on liquidity measures.

In the varimax interpretation of the loadings, each variable's factor assignment

is simplified and clarified even further. In the same vein as the Kaiser criterion, we determine a loading threshold of $\frac{1}{\sqrt{13}} \approx 0.277$. If a factor's variance is split across 13 variables, then its loading is considered above the threshold if it is greater than what an average variable might be expected to contribute. After discarding loadings below the threshold, the remaining significant loadings of original magnitude are presented in table 5.3. With the exception of LogD and OR, all variables are attributed to one specific factor. From the varimax rotation, almost every variable has a non-zero loading for only one factor. This suggests that for the most part, the information present across the mostly distinct groups of variables constituting each component is unique.

Finally, it is interesting to consider the extent to which PCA explains liquidity measures, and which liquidity measures perhaps might benefit from further analysis. In table 5.4, the cumulative variance of each original liquidity measure as explained by the five principal components is presented. From this, we can observe that the variance of the depth and trade count measures are very well explained with more than 75% of each variable's variance accounted for. The proportional realized spread and ILLIQ are the least-well explained variables, suggesting their variances are better explained with later components. This in itself perhaps indicates that these measures have more nuanced relationships with the other variables.

5.4 Principal Components Analysis of Citigroup's Liquidity Measures

After performing PCA on 2,485 observations of 13 liquidity measures of Citigroup, the Kaiser criterion indicated that there were four significant components, which together account for 74.58% of the total variance. Table 5.5 shows each significant component and the proportion of variance it accounts for. The second and third components explain roughly the same variance while there is larger difference between the first and second as well as the third and fourth. The elbow of the scree plot in Figure 5-3

	PC1	PC2	PC3	PC4	PC5	Cumulative Variance Explained
N	89.89%	0.54%	2.44%	0.21%	0.13%	93.21%
\$D	80.99%	1.08%	6.24%	2.76%	0.12%	91.18%
T	83.17%	1.25%	6.37%	0.20%	0.01%	90.99%
PES	0.13%	80.49%	4.27%	1.71%	0.00%	86.60%
M	0.01%	70.10%	10.80%	3.48%	1.08%	85.46%
LogD	64.88%	2.59%	12.39%	2.12%	0.02%	82.01%
FR	72.88%	0.39%	1.64%	3.74%	0.45%	79.10%
CL	1.13%	13.59%	54.53%	0.11%	6.08%	75.45%
LR1	3.20%	1.63%	0.56%	49.41%	9.26%	64.06%
PQS	0.05%	0.54%	4.07%	21.54%	37.03%	63.23%
OR	0.05%	0.11%	44.17%	10.14%	2.78%	57.25%
Illiq	7.14%	4.37%	2.25%	34.05%	4.59%	52.40%
PRS	0.18%	0.50%	1.80%	0.00%	40.69%	43.16%

Table 5.4: Percentage of cumulative variance of each of Bank of America's original liquidity measures as explained by the five significant PCA components.

appears to be at component four, which suggests retaining three principle components. However, the scree plot shows two steps in the graph, which perhaps makes the scree plot qualitative criteria less reliable as it does not strictly conform to the standard elbow shape. For this reason and because retaining three principal components would only account for 66.25% of the variance, we will consider the first four components to be significant.

	Standard deviation	Variance	Proportion of Variance	Cumulative Proportion
PC1	2.014125	4.056700	31.21%	31.21%
PC2	1.605291	2.576960	19.82%	51.03%
PC3	1.406803	1.979095	15.22%	66.25%
PC4	1.040773	1.083208	8.33%	74.58%

Table 5.5: The principal components of Citigroup using PCA on 13 liquidity measures and the Kaiser criterion to determine significance.

The loadings produced from the PCA can be found in figure 5-4 and table 7.6. The first principal component explains 31.21% of the variance of the original data. As with Bank of America, the first component is roughly an average of depth and trade count measures, namely turnover, dollar depth, log depth, trade count, and flow ratio.

The second component is responsible for 13.6% of the variance and mostly captures




Figure 5-3: The scree plot of Citigroup's principal components resulting from PCA on 13 liquidity measures.




Figure 5-4: The principal component loadings of Citigroup resulting from PCA on 13 liquidity measures.

the proportional effective spread, the Martin Index, and liquidity ratio one. The proportional effective spread describes the spread from mid at which trades occur, while the Martin Index measures the absolute change in price between consecutive trades squared divided by the turnover. Similarly, LRI indicates the average price

change of a transaction. Perhaps the commonality in these three measures arises because they are all generally describing price changes between transactions, whether it be proportional to the mid, the turnover, or the average time in between trades.

It is interesting that ILLIQ has a relatively small loading in this measure at -0.104 because it is a very similar measure to the Martin Index and liquidity ratio one. However, ILLIQ (and the proportional realized spread) are considering the price change from the beginning of the five-minute bar to the end of, instead of the change in price between consecutive trades that the Martin Index and liquidity ratio one measure.

The third component, describing 15.22% of the variance, shows the proportional quoted spread and composite liquidity contrasted against order ratio, the Martin Index, and liquidity ratio one. Composite liquidity being equal to the proportional quoted spread divided by the dollar depth may explain why their loadings are correlated and quite similar in value across all components.

For component five, the contrast between the proportional quoted spread and order ratio against the proportional realized spread and ILLIQ describe 8.33% of the variance. This contrast is counterintuitive as it suggests that when the proportional quoted spread and order ratio decrease (increase), the price changes become larger (smaller). Perhaps this can be rationalized by the quoted sizes decreasing overall, causing their difference to be smaller and the overall order ratio measure to be larger.

These interpretations of the components are supported by the varimax rotation of the loadings from the four significant principal components as seen in Table 5.7. The only noticeable difference is in component four where the loading on the proportional quoted spread is considered insignificant by our applied Kaiser criterion threshold of $\frac{1}{\sqrt{13}} \approx 0.277$ and the loading on log depth is considered significant. Perhaps this provides more insight into the meaning behind this fourth component. The negative correlation between log depth and order ratio does suggest that the order ratio is decreasing not necessarily because of increased liquidity but potentially instead because of an overall decrease in quote size.

Aside from log depth, each original variable is mapped to only one component. As

	PC1	PC2	PC3	PC4
PQS	-0.093	-0.288	0.478	0.394
PES	0.066	-0.567	-0.180	0.010
PRS	0.040	0.019	-0.031	-0.470
T	0.449	0.145	-0.122	0.108
\$D	0.454	0.025	0.183	0.012
LogD	0.386	-0.073	0.224	-0.174
CL	-0.050	-0.291	0.591	0.123
OR	-0.096	0.082	-0.328	0.633
N	0.468	0.086	-0.032	0.141
FR	0.392	0.097	0.041	0.242
M	0.061	-0.486	-0.299	-0.042
LR1	0.183	-0.462	-0.280	0.008
Illiq	0.085	-0.104	0.111	-0.293

Table 5.6: The principal component loadings of Citigroup using PCA on 13 liquidity measures and the Kaiser criterion to determine significance.

with Bank of America, this suggests that the information in these groups of variables identified by the components is unique and all provide insights into different and separate dimensions of liquidity.

	PC1	PC2	PC3	PC4
PQS			0.676	
PES		-0.59		
PRS				-0.408
LogD	0.344			-0.324
\$D	0.459			
N	0.49			
T	0.465			
CL			0.652	
M		-0.571		
LR1		-0.562		
FR	0.447			
OR				0.722
Illiq				-0.334

Table 5.7: A varimax rotation of Citigroup's first four principal component loadings resulting from principal components analysis on liquidity measures.

When examining the amount of correlation of each original variable as explained by the significant components, order ratio, the proportional realized spread, and ILLIQ are the three variables whose variances are least explained. As noted in Table 5.8, the

proportional realized spread and ILLIQ's variances are particularly underexplained at 24.87% and 17.46% respectively. Because both the proportional realized spread and ILLIQ use one price at the beginning of the five-minute bar and one price at the end of the five-minute bar in their calculations, the reliance on only two prices over the course of the five-minute bar could lead them to be more susceptible to noise, resulting in variances that can be less easily explained.

	PC1	PC2	PC3	PC4	Cumulative Variance Explained
CL	1.01%	21.83%	69.15%	1.64%	93.64%
N	88.89%	1.91%	0.20%	2.15%	93.15%
T	81.82%	5.42%	2.95%	1.26%	91.45%
PES	1.77%	82.88%	6.41%	0.01%	91.07%
\$D	83.65%	0.16%	6.63%	0.02%	90.46%
PQS	3.51%	21.38%	45.24%	16.82%	86.95%
LR1	13.59%	55.03%	15.52%	0.01%	84.15%
M	1.51%	60.89%	17.70%	0.19%	80.29%
LogD	60.47%	1.37%	9.93%	3.28%	75.06%
FR	62.36%	2.43%	0.33%	6.35%	71.47%
OR	3.74%	1.73%	21.30%	43.42%	70.19%
PRS	0.65%	0.09%	0.19%	23.94%	24.87%
Illiq	2.93%	2.79%	2.44%	9.30%	17.46%

Table 5.8: The variance of Citigroup's original variables explained by the significant principal components.

5.5 Principal Components Analysis of the Financial Sector ETF's Liquidity Measures

The principal component analysis of the Financial Sector ETF's overall liquidity measures encompasses 2,496 time periods containing 13 liquidity measure observations each. Five of those resulting principal components have eigenvalues greater than one, which is the Kaiser criterion for determining whether a principal component is significant or not. Together these five components account for 72.90% of the variance. The first component accounts for the most variance by over three times as much as the second component, the most of any first component out of all of the securities

analyzed. While the first component covers 34.57%, the remaining four components account for similar amounts of variance ranging from 11.22% to 8.18%. While the scree plot is shown in Figure 5-5, the location of the elbow in the scree plot is not entirely clear as cases could be made for its designation at either component two or component seven. Choosing only the first component to use in further analysis seems inadequate as it only describes a third of the variance; consequently, we consider the first five principal components, in between the values indicated by the scree plot, to be significant.

	Standard deviation	Variance	Proportion of Variance	Cumulative Proportion
PC1	2.119822	4.493645	34.57%	34.57%
PC2	1.207744	1.458645	11.22%	45.79%
PC3	1.138251	1.295616	9.97%	55.75%
PC4	1.079473	1.165262	8.96%	64.72%
PC5	1.031182	1.063337	8.18%	72.90%

Table 5.9: The principal components of the Financial Sector ETF using PCA on 13 liquidity measures and the Kaiser criterion to determine significance.

The loadings for the first five components are represented in Figure 5-6 and Table 5.10. The first component is an average of the same depth and trade count measures in other securities' first components (turnover, dollar depth, log depth, trade count, and flow ratio). The contrast between the proportional quoted spread and order ratio versus the proportional effective spread and composite liquidity synthesizes component two. Composite liquidity has the largest magnitude of the loadings while the proportional quoted spread, proportional effective spread, and order ratio have relatively similar loadings in magnitude. The contrast between composite liquidity and order ratio is one that has consistently appeared during the principal component analysis of these securities.

As discussed earlier, this is most likely a second order effect of both of the quoted size changing, rather than one increase or decreasing unevenly. This is further supported by the fact that the proportional quoted spread and composite liquidity are negatively correlated. Because the proportional quoted spread is the numerator of composite liquidity and dollar depth is the denominator, if the proportional quoted




Figure 5-5: The scree plot of the Financial Sector ETF's principal components resulting from PCA on 13 liquidity measures.

spreads and composite liquidity have different signs, it is an indication that this difference is mostly likely due to the denominator or depth component.

Component three has a variety of variables with mildly significant loadings. Generalizing, it shows the contrast of the proportional realized spread against the Martin Index, proportional quoted spread, proportional effective spread, liquidity ratio one and ILLIQ. Temporarily ignoring the negative loading of ILLIQ, this component might be a measure of the change in price of consecutive trades compared to the change in price over the five-minute bars. However, since ILLIQ does take into account the price return over five-minute bars, this might be furthered by noting that the large price changes experienced over five-minute bars that are not shown in consecutive trades are most likely accompanied by a large turnover. From this we

might infer that when the price of the Financial Sector ETF changes noticeably or trends, it takes a significant volume to do so.

If interpreted in this way, then the fourth component, which summarizes the contrast between the proportional realized spread and ILLIQ, is in some ways just an extension of the third component. This fourth component explains about a percent less of variance than the third component, accounting for 8.96% of the variance.

Finally, the fifth component covers the contrast between liquidity ratio one and the group of the proportional effective spread and order ratio. The component accounts for a similar amount of variance as component four at 8.18%.




Figure 5-6: The principal component loadings of the Financial Sector ETF resulting from PCA on 13 liquidity measures.

Table 5.11 shows the varimax rotation of the first five principal components' loadings. The interpretations from the original and rotated loadings are consistent for the first, second, and fourth principal components. However, the interpretations vary slightly for the third and fifth components. The third is represented as an average of the proportional effective spread and the Martin Index, or rather some measure of the price change in between consecutive transactions. And the fifth can be summarized as an average of the proportional quoted spread, the Martin Index, and liquidity ratio one.

	PC1	PC2	PC3	PC4	PC5
PQS	-0.013	0.391	-0.379	0.166	0.188
PES	0.122	-0.333	-0.309	0.402	-0.426
PRS	-0.003	0.191	0.319	0.576	0.162
T	0.415	0.203	0.057	-0.026	-0.079
\$D	0.440	-0.138	0.135	0.039	-0.075
LogD	0.399	-0.229	-0.011	-0.068	0.244
CL	0.006	-0.616	-0.033	-0.016	0.069
OR	0.051	0.378	-0.329	-0.015	-0.478
N	0.452	0.091	0.093	0.028	-0.038
FR	0.396	0.084	0.194	0.074	-0.240
M	0.028	-0.179	-0.512	0.428	0.091
LR1	0.277	0.123	-0.303	-0.016	0.613
IlliQ	0.137	-0.037	-0.359	-0.531	-0.090

Table 5.10: The principal component loadings of the Financial Sector ETF using PCA on 13 liquidity measures and the Kaiser criterion to determine significance.

	PC1	PC2	PC3	PC4	PC5
PQS		0.386			0.418
PES			-0.719		
PRS				0.699	
LogD	0.308	-0.306			
\$D	0.46				
N	0.459				
T	0.427				
CL		-0.537			
M			-0.598		0.318
LR1					0.741
FR	0.48				
OR		0.627			
IlliQ				-0.647	

Table 5.11: A varimax rotation of the Financial Sector ETF's first five principal component loadings resulting from principal components analysis on liquidity measures.

The ways in which the variance accounted for is distributed across the original variables is quite different than the stocks. The proportional realized spread and ILLIQ are no longer the least explained variables as composite liquidity and the proportional quoted spread are now, at 56.06% and 46.97% respectively. However, 55.37% of the variance of composite liquidity is covered in the second principal component. So although a large variance of composite liquidity is not captured by all of

the components, this does suggest that component two is able to capture a meaningful dimension of composite liquidity.

The opposite can be said for components such as the Martin Index, order ratio, the proportional effective spread, and even liquidity ratio one. While a large amount of variance of these measures is accounted for by the first five principal components, the variances are divided relatively equally across some or even all of the components. This can make the contributions and interactions of these variables more difficult to interpret and draw decisive conclusions from. Bank of America and Citigroup do not demonstrate this same dynamic concerning the proportional effective spread; in fact, that is one variable that has most of its variance allocation to one component for both stocks.

	PC1	PC2	PC3	PC4	PC5	Cumulative Variance Explained
N	91.84%	1.21%	1.12%	0.09%	0.15%	94.42%
\$D	87.03%	2.78%	2.36%	0.18%	0.60%	92.95%
LR1	34.49%	2.21%	11.90%	0.03%	39.97%	88.60%
LogD	71.57%	7.65%	0.02%	0.54%	6.33%	86.11%
T	77.42%	6.01%	0.42%	0.08%	0.66%	84.60%
FR	70.50%	1.03%	4.88%	0.64%	6.13%	83.17%
PES	6.69%	16.18%	12.38%	18.84%	19.30%	73.39%
M	0.35%	4.68%	33.98%	21.35%	0.88%	61.24%
OR	1.17%	20.85%	14.03%	0.03%	24.31%	60.38%
PRS	0.00%	5.32%	13.19%	38.68%	2.79%	59.98%
Illiq	8.44%	0.20%	16.70%	32.87%	0.86%	59.07%
CL	0.02%	55.37%	0.14%	0.03%	0.51%	56.06%
PQS	0.08%	22.31%	18.62%	3.21%	3.76%	47.97%

Table 5.12: The variance of the Financial Sector ETF's original variables explained by its five significant principal components.

Chapter 6

Time Series Analysis of Principal Component Scores of Liquidity Measures

Because principal components analysis is an analysis of spatial variance, time is not explicitly captured in the analysis of liquidity measures in Chapter 5. Even so, the results from the analysis still suggest there is a temporal nature to the component scores. More specifically, the score plots of the components show clear liquidity trends evolving over time, potential autocorrelation, and possibly even cross-correlation between scores. Verifying and quantifying the autocorrelation and cross-correlation present in these scores can be useful in selecting a model that aims to predict future values and thus future liquidity.

In this case of our data, the autocorrelation present suggests the use of a vector autoregressive (VAR) model. This model results in estimates of the influence of lagged liquidity measures on the liquidity in time t , and is, therefore, an approach that may lead to promising results in predicting liquidity. While we apply VAR to the scores of the principal components, to better predict changes in liquidity, it would also be appropriate to use the first differences of the scores as input data to the VAR model.

6.1 Autocorrelation and Cross-correlation in Liquidity Measures

Autocorrelation is the correlation of a vector with a lagged version of itself. It is useful for identifying time dependencies in data, which aid in the modeling of time series in such a way that enables the prediction of future values from current and past values. The presence of autocorrelations is a condition for applying a VAR model to a time series.

Autocorrelation is calculated in the same way that the correlation between two separate times series is, however, in this case, one time series is a lagged version of itself. More formally, given a vector of score values for a component, Y , where each row y_1, y_2, \dots, y_n is a linear combination of original variables x_1, x_2, \dots, x_p , the lag k autocorrelation, r_k , is defined as:

$$r_k = \frac{\sum_{i=k+1}^N (Y_i - \bar{Y})(Y_{i-k} - \bar{Y})}{\sum_{i=1}^n (Y_i - \bar{Y})^2}$$

Although the time of the scores Y is not used in the calculation of autocorrelation, it is assumed that the observations of X are eqispaced. In the context of this analysis, the observations are spaced at five-minute intervals.¹

Autocorrelation is a correlation coefficient that takes a value between -1 and measures the correlation between two values of the same variable at times t and $t - k$. When using autocorrelation to identify an appropriate times series model, autocorrelations for multiple lags are plotted sequentially. From these plots, statistical significance can determine if the series exhibits autocorrelation. Cross-correlation closely follows autocorrelation with the exception that it is the correlation between a variable and a lagged version of a different variable. In the case of principal component scores, it is the correlation between one component's scores and a different compo-

¹During each day, the bars are equi-spaced; however, overnight gaps and weekends introduce irregular spacing into the time series, specifically between the last measure of a day and the first measure of the following day. This irregular spacing is not adjusted for in any way, and consequently, the VAR models of these time series are most suitable in the context of daily weekday liquidity, rather than liquidity spanning multiple days.

ment's scores at lag k .

In order to determine the number of lags to include in the model, or rather what the order should be, partial autocorrelation is often first qualitatively examined. Partial autocorrelation is the amount of correlation between a variable and a lag of another variable or itself that is not explained by correlations at all lower-order lags. For instance, the autocorrelation of a time series Y at lag 1 is the coefficient of correlation between Y_t and Y_{t-1} , which is presumably also the correlation between all other pairs of consecutive observations in the series. But if Y_t is correlated with Y_{t-1} , and Y_{t-1} is equally correlated with Y_{t-2} , then one would expect the correlation between Y_t and Y_{t-2} to be the square of the lag-1 correlation. Thus, the correlation at lag 1 propagates to lag 2 and presumably to higher-order lags. The partial autocorrelation at lag 2 is therefore the difference between the actual correlation at higher-order lag and the expected correlation due to the propagation of correlation at all lower-order lags. The partial autocorrelation is useful for identifying the order of the autoregressive model that can be used to fit the time series. The partial autocorrelation is calculated as such:

$$\alpha(\rho_k) = \begin{cases} r_1 & \text{for } k = 1 \\ \frac{r_k - \sum_{j=1}^{k-1} \rho_{k-1,j} \cdot r_{k-j}}{1 - \sum_{j=1}^{k-1} \rho_{k-1,j} \cdot r_{k-j}} & \text{for } k > 1 \end{cases}$$

and

$$\rho_{k,j} = \rho_{k-1,j} - \rho_k \cdot \rho_{k-1,k-j}$$

6.2 Vector Autoregressive Model

A vector autoregressive (VAR) model leverages linear interdependencies among a multivariate time series to create a model that enables the forecasting of future values using past and current values. The m -dimensional multivariate time series X_t follows the VAR(p) model with auto-regressive order p if

$$x_t = C + \phi_1 x_{t-1} + \phi_2 x_{t-2} + \cdots + \phi_p x_{t-p} + \eta_t$$

where $C = (c_1, c_2, \dots, c_m)'$ is an m -vector of constants, $\phi_1, \phi_2, \dots, \phi_p$ are $(m \times m)$ matrices of coefficients, and η_t is white noise. For this analysis, the matrix X_t consists of the scores of the significant principal components.

In order to determine the appropriate order (p) of the model, the Schwarz criterion (SC), also known as the Bayesian information criterion, can be applied. It is defined as:

$$SC = \ln(n)k - 2\ln(\hat{L})$$

where n is the number of observations, k is the number of the parameters estimated by the model, and \hat{L} is the maximized value of the likelihood function with k parameters estimated from n observations.

6.3 Bank of America Time Series Analysis

As the cumulative sums of component scores can provide insight into the univariate time series structure of each component's scores, Figure 6-1 displays the the cumulative sum of Bank of America's component scores over time. Despite the fact that the principal components of Bank of America seemed to account for mostly disjoint sets of measures, the cumulative sum of the component scores in Figure 6-1 suggest that some of these components move in tandem.

For instance, the score graphs of components one and three look similar in shape (both start off in a smaller range, increase to a global maximum, decrease to a global minimum, and then end somewhere in between the minimum and maximum). However, component one lags component three by about two business days. Component three mostly covers quoted prices and sizes while component one concerns statistics of realized trade volumes and counts. Perhaps component one lagging component three conveys market participants responding to the changing market maker prices and conditions. Components four and five also share a similar shape in terms of magnitude. Component five lags component four up until the end of February where the lag reverses. Overall, the score plots show Bank of America's liquidity beginning

at moderate levels, increasing, decreasing, and then returning to moderate levels over the time period of the analysis.




Figure 6-1: The cumulative sum of Bank of America’s principal components scores over time.

6.3.1 Time Series Structure

The autocorrelations of each of Bank of America’s principal component’s scores and their cross-correlations are shown in Figure 6-2. Because a zero-lag transformation of a measure is equivalent to the original measure itself, the zero lag has a correlation coefficient of one for all of these component autocorrelation plots. However, significance beyond the zero lag indicates autocorrelation.

Out of all of the components, component one demonstrates the most significant autocorrelation. Figure presents the autocorrelation and partial correlation collectively. Component one has positive autocorrelation which begins on the first lag and

decays geometrically until it becomes statistically insignificant after a one-hour window (the twelfth lag). As described in Chapter 5.3, the first principal component largely captures the weighted average of the volume, depth, and trade count measures (turnover, dollar depth, logarithmic depth, trade count, and flow ratio). So this suggests that these previous measures impact the current measures for a rolling one-hour window and are positively correlated.

In addition to the positive autocorrelation with lags one through twelve, component one also demonstrates a slight negative correlation around the 25-bar lag. While a 25-bar window may initially seem arbitrary, 25 five-minute bars is roughly equivalent to a third of the trading day. Given that component one describes variables that have demonstrated intraday seasonality (as shown in Chapter 4), perhaps the 25-bar lag is encompassing the daily U-shape pattern that the underlying variables of component one exhibit.

The partial autocorrelation of component one in Figure 6-3 shows the first four lags as significant, thus suggesting a fourth order autoregressive model could be applied to this time series. The effectiveness of this could potentially be impacted by the ratio of large to small trades during a time period. For instance, large block trades are more likely to be split across many small trades over a period longer than twenty minutes, in which case, perhaps a longer window may be more effective in some scenarios.

The second component has statistically significant autocorrelation when a one-bar lag is applied as indicated in Figure 6-4. While lag 24 is also slightly significant, it is most likely the result of noise, considering lags two through 23 are not. However, the magnitude of the lags does seem to present in a pattern that could indicate seasonality as it goes through cycles of increasing and decreasing in magnitude. The partial autocorrelation, also in Figure 6-4 shows that only the first lag is significant, which suggests a first-order autoregression model may be appropriate.

Given that this second factor describes the variance of price changes between trades (the Martin Index) and the spread from mid at which a trade occurs (the proportional effective spread), a lag of one suggests that previous price changes and previous effective spreads from the past five minutes have forecasting power in the

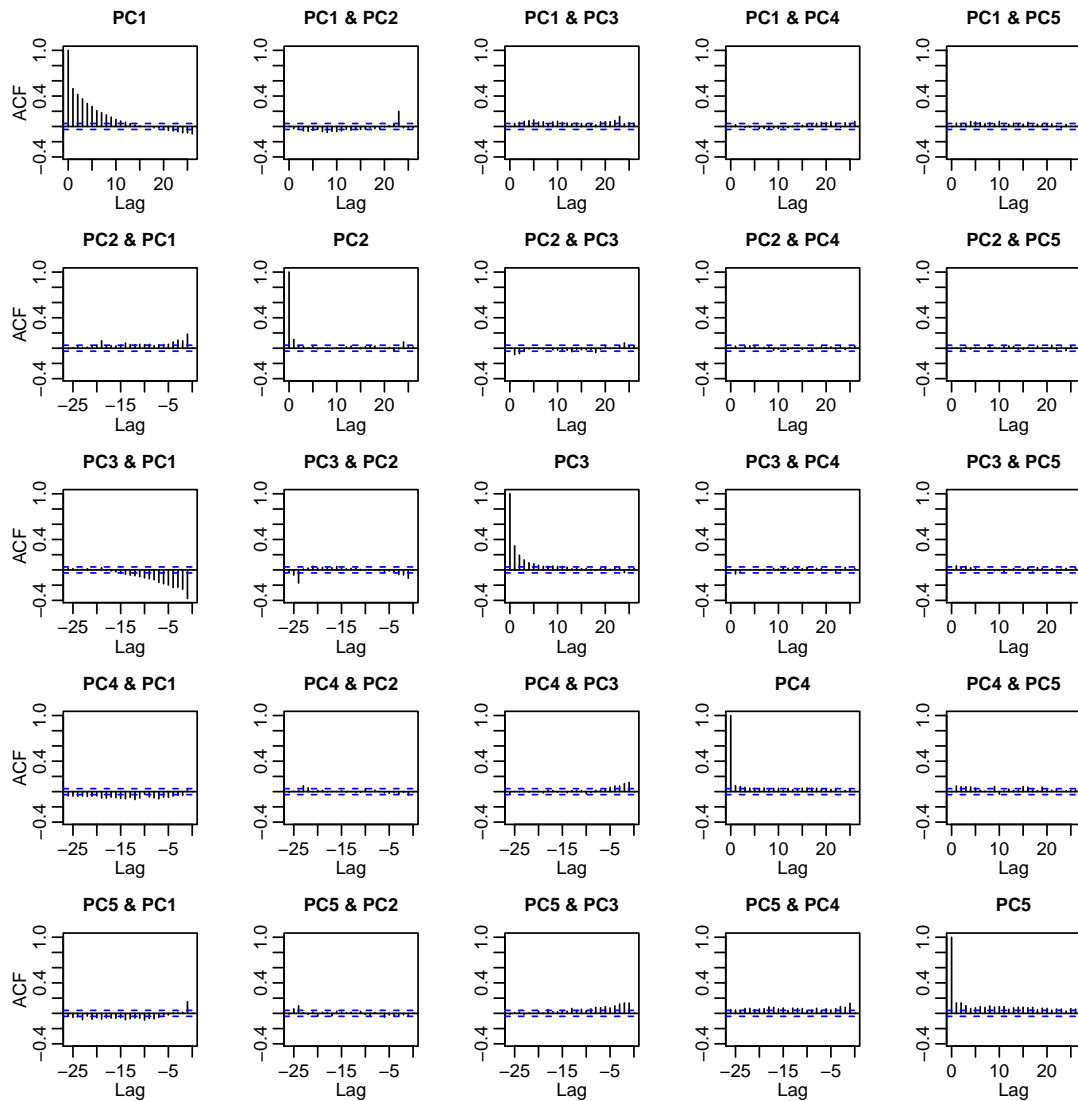


Figure 6-2: The autocorrelation and cross-correlation of Bank of America’s principal component scores.

current five minutes. The absence of significant correlation to other lags indicates that the Martin Index and the proportional effective spread are short-term measures whose forecasting abilities range from five to ten minutes before becoming stale. The short-term dependence of price change measures and spread measures suggests that while there may have been an overarching liquidity theme for Bank of America over the course of the time period of the analysis, within this trend, there is variance in

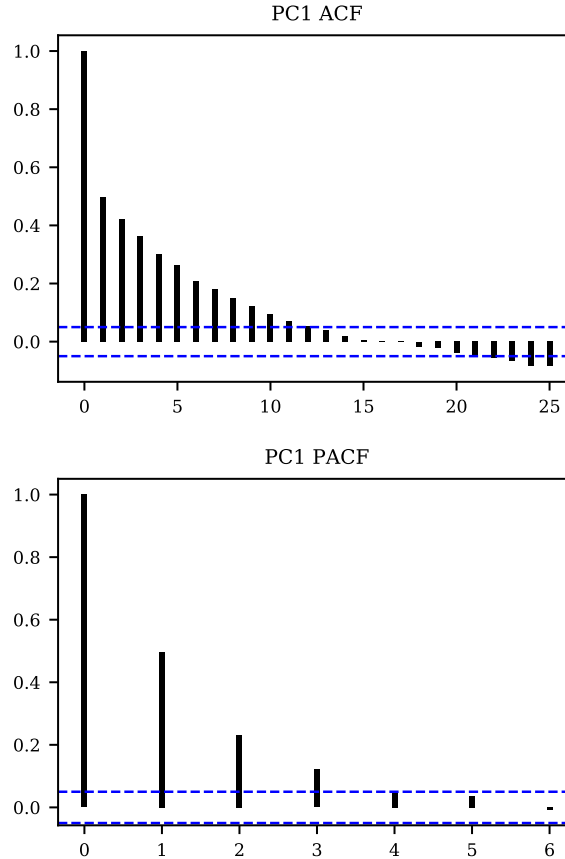


Figure 6-3: The autocorrelation and partial autocorrelation of Bank of America’s first principal component’s scores.

the variables on an intraday and daily basis.

The third principal component contrasts order ratio and composite liquidity, interpreted in Chapter 5 as the difference between how the bid and ask quote sizes are adjusted compared to how their spread is adjusted. Shown in Figure 6-5, component three’s autocorrelation is significant starting at the first lag, and it geometrically decreases until it becomes statistically insignificant after the sixth lag. This indicates a 30-minute window for which the previous differences between how the bid and ask quote sizes are adjusted compared to how their spread is adjusted are relevant for the current five-minute bar.

While the interpretation of the third component is perhaps less conceptually tractable than other components, given that it seemed to exhibit cross-correlation




Figure 6-4: The autocorrelation and partial autocorrelation of Bank of America’s second principal component’s scores.

in Figure 6-1 with component one and that component one showed instances of autocorrelation, it stands to reason that component three might likely exhibit time dependencies. Its partial autocorrelation plot indicates that the first two lags are significant, which suggests a second-order autoregression model.

Although the time structures of the first three components are meaningful, the fourth component has only slightly significant autocorrelation. As shown in Figure 6-6, the correlations of the first two lags are greater than 0.05, which indicates significance at a 95% confidence level. The partial autocorrelation suggests a first-order model.

Because the fourth component captures a weighted average of liquidity ratio 1, which is the average price change of a transaction, and ILLIQ, which is defined as the




Figure 6-5: The autocorrelation and partial autocorrelation of Bank of America's third principal component's scores.

return over a five minute bar divided by the turnover during that bar, perhaps, it does not demonstrate very meaningful time series structure because of the wide-ranging aspects of liquidity (price changes of consecutive trades, trade count, price return over a longer term (the five-minute bar), and turnover) that it covers. Additionally, the lack of substantial five-minute bar time series structure does not exclude the possibility that this exists in the fourth component; it just may be present in bars that are smaller or larger than five minutes.

Lastly, the fifth component, which describes the contrast between the proportional quoted spread and the proportional realized spread, presents autocorrelation. Figure 6-7 shows there is significant correlation for a number of lags, and that correlation rises and falls as the overall magnitude decreases. The way in which these autocorrelations




Figure 6-6: The autocorrelation and partial autocorrelation of Bank of America’s fourth principal component’s scores.

oscillate suggests indicates perhaps some degree of seasonality in component five.

As postulated in Chapter 5.3, perhaps this seasonality is an indication of market makers using the previous bar’s price changes to adjust their quotes during the current bar. Their adjustment in the current bar could then in turn cause the realized spread’s behavior to change, creating almost a self-perpetuating cycle of sorts. Nevertheless, lags one through three have the largest correlations and partial correlations where a second-order autoregression model may be appropriate.

While all components exhibited some degree of autocorrelation at the 95% confidence level, the same cannot be said for the cross-correlation between components. Because components one and three each individually showed the strongest autocorrelation trends among components, perhaps it is unsurprising that their cross-




Figure 6-7: The autocorrelation and partial autocorrelation of Bank of America’s fifth principal component’s scores.

correlation is the most statistically significant among all combinations of components. Figure 6-2 shows that component one lags component three by one five-minute bar with a correlation coefficient of -0.38. From there, the strength of the correlation geometrically decreases until it becomes insignificant after lag 15.

Because the third component contains largely market maker quoted spreads and quoted depths, this suggests that volume traded, number of trades, and other volume/depth measures are reactionary in part to these quotes. The partial autocorrelation between components three and one shows the first lag is significant, which indicates a first-order model. Although there are other small instances of cross-correlation between components that can demonstrate time dependencies, these dynamics are perhaps better suited for analysis using the results of the VAR model

where the associations become more clear.

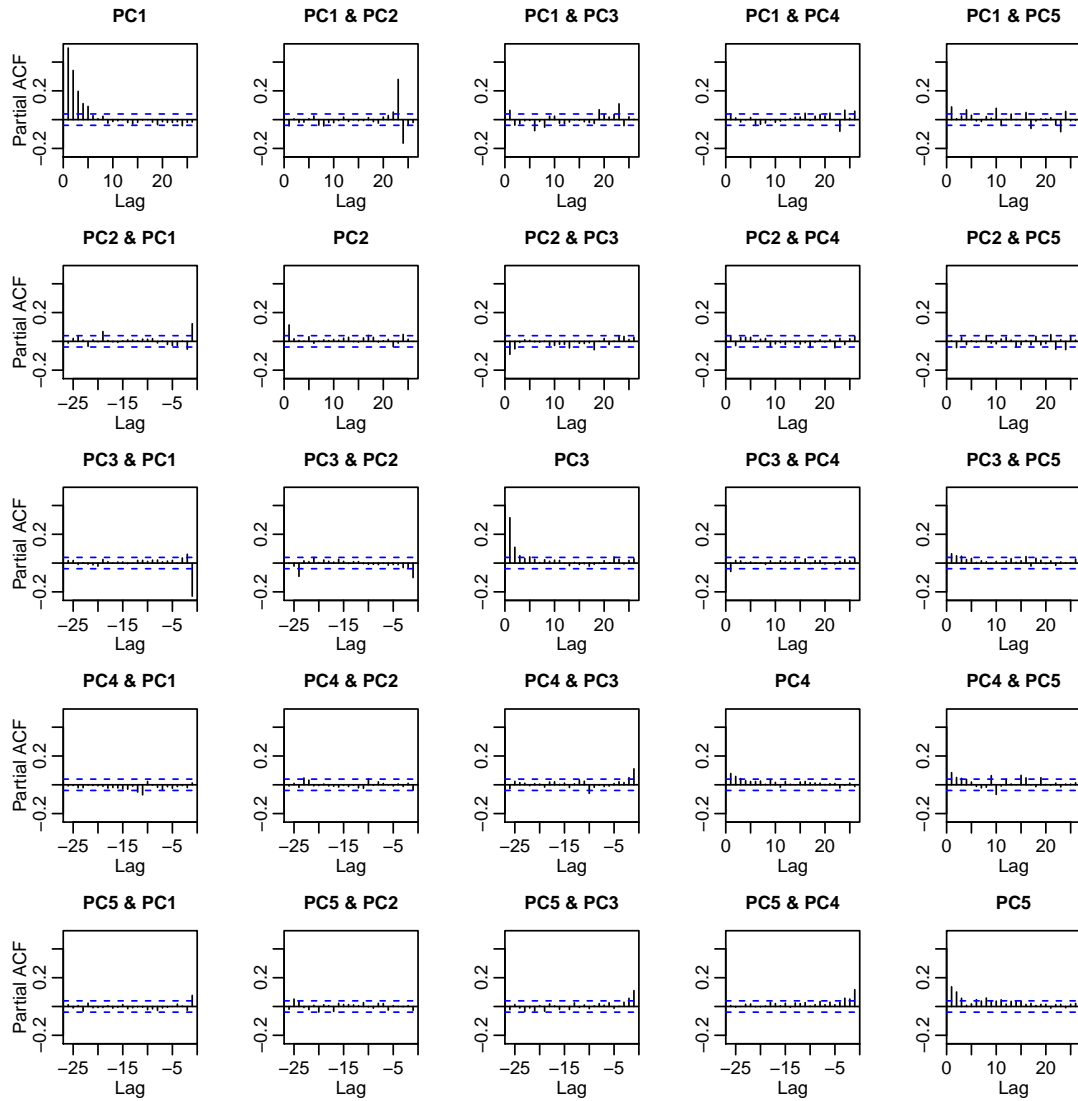


Figure 6-8: The partial autocorrelation of Bank of America's principal component scores.

6.3.2 Vector Autoregressive Model

Although from qualitative observation of the partial autocorrelations plots for all components and their combinations (Figure 6-8), it is unclear which order VAR model should be used, we use the SC to analytically determine this. Since the SC reaches its

maximum value at two, a second order VAR model is applied to these five principal components. The results of the model are contained in Tables 6.1 through 6.5.

Table 6.1 shows the results for the first component where component one is statistically significant on a 0.10% level on the first and second lags of itself and the first lag of components two and three. Component one additionally has slight dependence on the first lag of component four (at a level of at least 5%). This suggests that volume traded and the frequency of trades is largely dependent on its previous values in the last ten minutes and price changes and quoted spreads and depths of the previous five minutes.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.3358	0.0206	16.29	0.0000 ***
PC2.11	-0.1601	0.0264	-6.06	0.0000 ***
PC3.11	0.3071	0.0333	9.22	0.0000 ***
PC4.11	-0.1047	0.0442	-2.37	0.0180 *
PC5.11	0.0519	0.0389	1.34	0.1817
PC1.12	0.3491	0.0231	15.12	0.0000 ***
PC2.12	-0.0050	0.0258	-0.19	0.8466
PC3.12	-0.0240	0.0299	-0.80	0.4230
PC4.12	-0.0782	0.0445	-1.76	0.0786
PC5.12	0.0544	0.0375	1.45	0.1475
const	0.0106	0.0332	0.32	0.7485

Table 6.1: The first principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

Table 6.2 displays the summary statistics for the VAR model for component two. It shows very similar results to those of component one. At the 0.10% level, component two is influenced by the lag one value of itself and components one through three. Interestingly, component two has a second order dependence on component at a 1% level and on component three at 5% level, but not on itself. Additionally, it depends on the first lag of component four slightly (5% level).

Component two describes price changes of consecutive trades and the price from mid at which a trade occurs. These results indicate that price changes and customer execution costs largely rely on previous price change and effective spread in the last

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.1513	0.0161	9.41	0.0000 ***
PC2.11	0.1274	0.0206	6.18	0.0000 ***
PC3.11	-0.1090	0.0260	-4.20	0.0000 ***
PC4.11	0.0729	0.0345	2.11	0.0346 *
PC5.11	0.0279	0.0303	0.92	0.3571
PC1.12	-0.0561	0.0180	-3.11	0.0019 **
PC2.12	0.0197	0.0201	0.98	0.3281
PC3.12	-0.0545	0.0233	-2.34	0.0196 *
PC4.12	-0.0139	0.0347	-0.40	0.6883
PC5.12	-0.0513	0.0293	-1.75	0.0798
const	-0.0004	0.0259	-0.01	0.9881

Table 6.2: The second principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

five minutes and quoted spreads and sizes, volume traded, and number of trades in the previous 10 minutes. Because there is no second-order dependence on itself, volume and depth measures seem to be more influential for a longer period of time than the price changes and effective spread. In other words, the volume being traded, the number of trades occurring, and the subsequent quoted prices and depths are driving factors in what the price changes will be.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	-0.2574	0.0129	-20.02	0.0000 ***
PC2.11	-0.1156	0.0165	-7.01	0.0000 ***
PC3.11	0.3213	0.0208	15.46	0.0000 ***
PC4.11	-0.1687	0.0276	-6.12	0.0000 ***
PC5.11	0.0615	0.0242	2.54	0.0113 *
PC1.12	0.0614	0.0144	4.26	0.0000 ***
PC2.12	-0.0371	0.0161	-2.31	0.0210 *
PC3.12	0.1177	0.0187	6.31	0.0000 ***
PC4.12	-0.0256	0.0277	-0.92	0.3559
PC5.12	0.0733	0.0234	3.13	0.0018 **
const	0.0019	0.0207	0.09	0.9272

Table 6.3: The third principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

Component three's coefficients and errors are displayed in Table 6.3. Component three depends strongly on a number of components at a number of lags. At the 0.01% level, there is a first and second order dependence on itself and component one and just a first order dependence on components two and four. Furthermore, there is a first and second order dependence (at the 5% and 1% levels respectively) on component five.

Component three is mostly determined by the contrast between composite liquidity and order ratio; in other words, the difference between the quoted spread and the quoted depth imbalance. Thus this is a component that is largely the result of only market makers' input. To a certain extent, as discussed in Chapter 5.3, it even indicates market makers' speculation if we consider this contrast between quoted spread and depth to be an indication of such. The high significance on such a large number of components suggests that market makers are using all available information from the previous five minutes in a somewhat systematic fashion to determine what their current quotes will be. Furthermore, it seems that spreads and quoted depth depend on longer windows of multiple measures hence the strong second-order dependence.

Component four's results are contained in Table 6.4, and they indicate significant first and second order dependencies on itself and component three. The average price change of a transaction and the 5-minute bar return over the bar's turnover dominate the fourth component. The existence of a dependency between quoted spread and depth and price change is not surprising. In fact, smaller quoted depth on one side or both can cause significant price changes when transactions of larger size are forced to walk down the order book and execute at a variety of prices. It is interesting that component four has a greater dependency on its second lag rather than its first. Perhaps this indicates mean-reversion tendencies of average transaction cost and ILLIQ over a 15-minute time period.

Finally component five relies heavily on the the first and second lags of all first-order components besides component two. Of the second-order components, the fifth is influenced by itself at the 0.01% level, component three at the 0.01% level, and components one and four at the 5% level. The fifth component captures the con-

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.0205	0.0139	1.47	0.1408
PC2.11	-0.0291	0.0179	-1.63	0.1034
PC3.11	0.0862	0.0225	3.83	0.0001 ***
PC4.11	0.0774	0.0299	2.59	0.0097 **
PC5.11	0.0606	0.0263	2.31	0.0212 *
PC1.12	-0.0237	0.0156	-1.52	0.1292
PC2.12	0.0161	0.0174	0.93	0.3546
PC3.12	0.0424	0.0202	2.10	0.0360 *
PC4.12	0.1033	0.0301	3.43	0.0006 ***
PC5.12	0.0267	0.0254	1.05	0.2937
const	-0.0018	0.0225	-0.08	0.9346

Table 6.4: The fourth principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

trast between the proportional quoted spread and the proportional realized spread, or rather the difference between liquidity risk priced in the market (as represented by the proportional quoted spread) and liquidity risk that is realized after five minutes. It is not surprising that a measure which describes all aspects of the trading process (from market makers' quotes to realized price movement after five-minutes) relies on components that generally isolate specific aspects of this process, the highly significant first and second-order dependence that component five has on itself is perhaps unexpected. This suggests that if there is a mismatch between risk being priced in for immediate execution (the proportional quoted spread) and the realized price movement after five minutes (the proportional realized spread), then this mismatch perpetuates forward.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.0843	0.0120	7.05	0.0000 ***
PC2.11	-0.0162	0.0153	-1.05	0.2919
PC3.11	0.0747	0.0193	3.87	0.0001 ***
PC4.11	0.1153	0.0256	4.50	0.0000 ***
PC5.11	0.1061	0.0225	4.71	0.0000 ***
PC1.12	-0.0264	0.0134	-1.97	0.0493 *
PC2.12	0.0077	0.0149	0.51	0.6080
PC3.12	0.0546	0.0174	3.15	0.0017 **
PC4.12	0.0635	0.0258	2.46	0.0139 *
PC5.12	0.0947	0.0218	4.35	0.0000 ***
const	0.0009	0.0193	0.04	0.9644

Table 6.5: The fifth principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

6.4 Citigroup Time Series Analysis

As described in Chapter 5.4, there are four principal components of Citigroup's overall liquidity statistics that are considered significant, which account for almost 75% of the variance of the data. The first component is responsible for 31.21% of this variance and is a volume and depth measure, which can be approximately described as an average of the trade count, turnover, dollar depth, log depth, and flow ratio. The second component is an average of the proportional effective spread, the Martin Index, and liquidity ratio one, and is a measure that generally captures price changes between trades, the average price change of a trade, and the spread from mid at which a trade occurs. The third component is an average of the proportional quoted spread and composite liquidity, and the fourth component contrasts the proportional realized spread, log depth, and ILLIQ against the order ratio. This fourth measure might be thought of as the difference between price returns over five minutes (either adjusted by turnover or the mid-price of the stock) and the quoted depth imbalance.

The scores of these four significant components are in Figure 6-9, and while each component is largely determined by a set of variables that is disjoint from the variables that describe the other components, the degree to which these components seem to




Figure 6-9: The cumulative sum of Citigroup’s principal components scores over time.

contain separate information is even more evident from the the cumulative sums of its scores over time. If we consider the first principal component to be an indicator of liquidity across depth and volume measures, then over the time period examined, liquidity seems to decrease, reaching a global minimum on February 28, 2019, and then increase for the remainder of the period, peaking on March 18, 2019. Principal component one also visually exhibits the largest range out of all components.

Principal components two and three appear bi-modal and also support the notion of liquidity decreasing and then increasing by the end of the analysis period. Although the scores of component two increase and then decrease, a larger component score corresponds to decreased liquidity because the loadings are negative on the main variables that describe component two. While component two spans a larger range than component three does, the magnitude of the trough between the two peaks

of component three is much larger than it is for component one. Since the trough between the two peaks of component three corresponds to the timing of the global minimum of component one, perhaps these measures are interrelated; this potential cross-correlation between component one and three might explain why it experiences more variation between peaks than component two.

The scores for component four also indicate a trend of declining liquidity over the period, which then reverts to its starting level of liquidity at the end of the period. While component four displays this general trend, it is not quite clear as the scores of the other components exhibit it. Perhaps, component four is a measure that is more short-term, and while it captures the same overall trend, its intraday variation is more of the dominant variance.

6.4.1 Time Series Structure

Given these qualitative long-term observations of the scores' time structures, the autocorrelations and cross-correlations of each component's scores in Figure 6-10 shows the autocorrelations and cross-correlations and Figure 6-11 shows the partial correlations of the principal component scores. From this, we can see that all principal components demonstrate statistically significant autocorrelation. Furthermore, in some instances, cross-correlations between components are statistically significant, indicating time interdependencies between the component scores.

Principal component one perhaps has the most significant autocorrelation of all the components. The first 11 lags are statistically significant, which suggests that volume, depth, and trade count measures from the past 55 minutes are relevant to the current time bar's component one score. Component two depends on a much shorter period of time as only its first lag is significant. Because the measures that component two includes are price change and spread measures, it stands to reason that measures that are inherently short-term would not have a large lag effect. Meanwhile, the third and fourth components both exhibit significant autocorrelation for multiple lags. Perhaps the structure of these is somewhat surprising as these components also include price change measures but exhibit much more time series structure. Perhaps

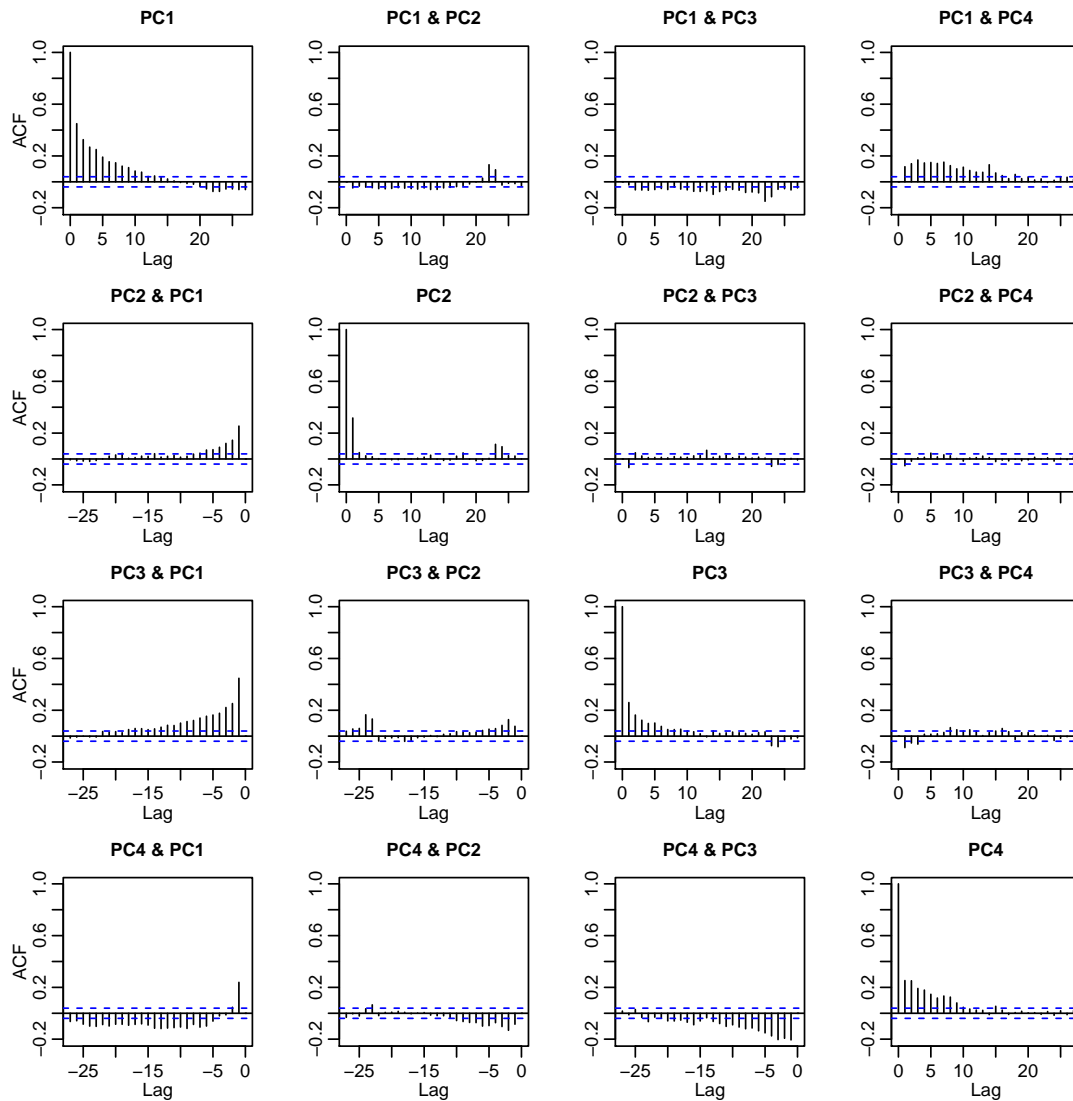


Figure 6-10: The autocorrelation and cross-correlation of Citigroup's principal components scores over time.

the depth element in component three, composite liquidity, and the log depth that factors into component four may be at least partly responsible for the autocorrelation.

Cross-correlation is prevalent in many of these combinations of components. For the most part, component one seems to lag behind components two, three, and four. The lag is positively correlated for components two and three, but negatively so for component four. Additionally, the lag coefficient is stronger on more recent measures

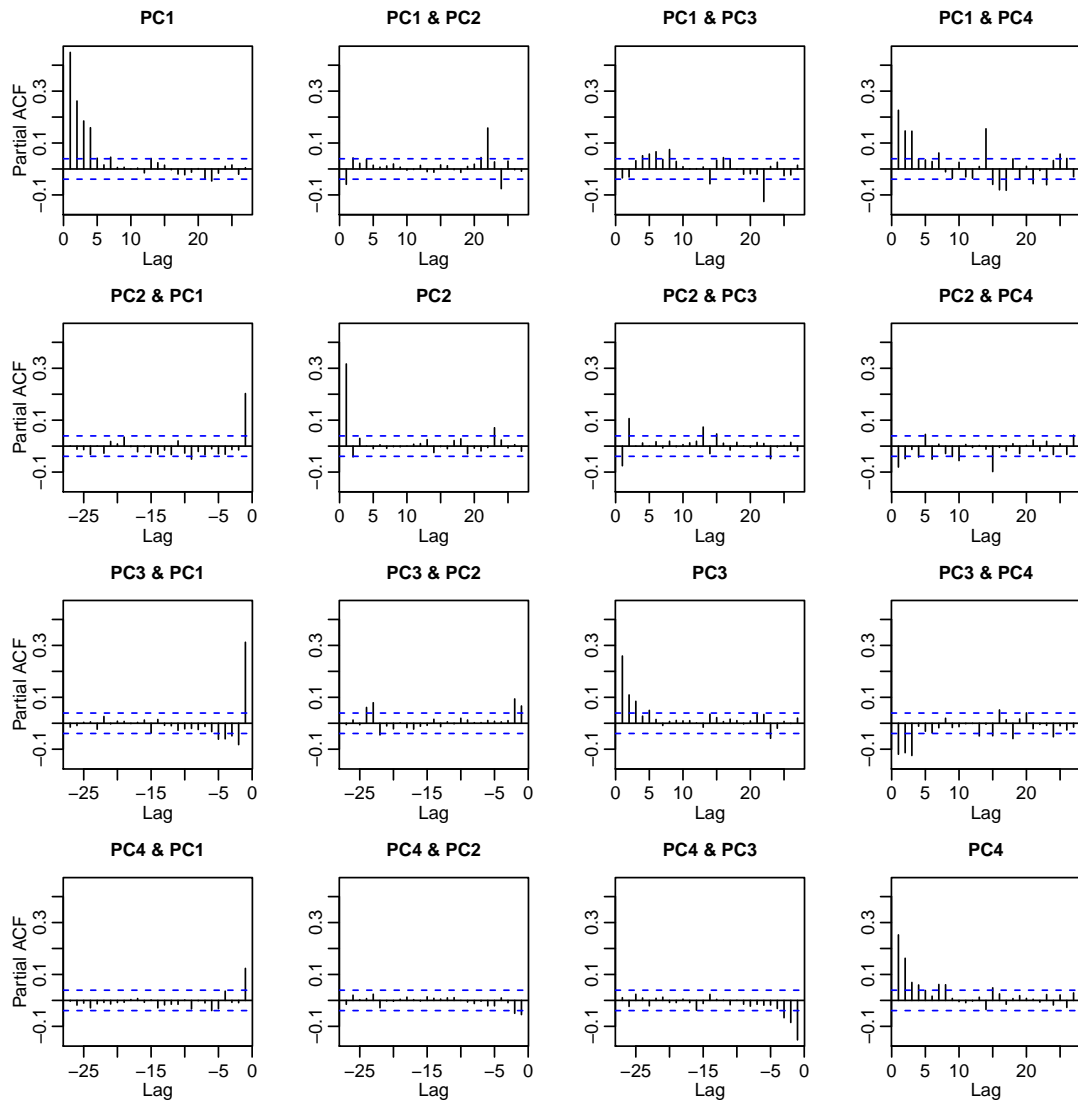


Figure 6-11: The partial autocorrelation and partial cross-correlation of Citigroup’s principal components scores over time.

when component one lags components two and three, but the significance of the lags for component four preceding component one is more dispersed.

The partial autocorrelations in Figure 6-11 suggest varying degrees of order for the VAR model to be fit to this data. While the partial autocorrelation for many of the components suggests a first order model, for some components such as the first, third, and fourth components, a third or fourth order model is indicated from

the degree of significant lags. When the SC is applied to the data, it indicates that VAR(3) is appropriate for this multivariate time series. After fitting a VAR(3) model to the scores, we see the many interdependencies between time and components demonstrated.

6.4.2 Vector Autoregressive Model

The results for the first principal component are shown in Table 6.6. The first component has a highly significant third order dependence on itself. Furthermore, the first lag of every component but the fourth and then the second lag of the first component are highly significant. Oddly, the first and second lags of component four are not statistically significant, but its third lag is highly significant, which indicates that the value of the fourth component's score 15 minutes before the current bar is meaningful. The first component is a volume, depth, and trade count measure, while the fourth component describes the contrast of price returns over five minutes against the order ratio, or rather an indicator of the buy and sell quoted size imbalance. The order

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.2884	0.0212	13.62	0.0000 ***
PC2.11	-0.1568	0.0246	-6.36	0.0000 ***
PC3.11	-0.1824	0.0315	-5.79	0.0000 ***
PC4.11	0.0372	0.0382	0.97	0.3305
PC1.12	0.1891	0.0275	6.89	0.0000 ***
PC2.12	-0.0205	0.0261	-0.79	0.4318
PC3.12	-0.1593	0.0323	-4.93	0.0000 ***
PC4.12	0.0504	0.0381	1.32	0.1860
PC1.13	0.1863	0.0259	7.21	0.0000 ***
PC2.13	0.0227	0.0233	0.97	0.3309
PC3.13	0.0326	0.0268	1.22	0.2236
PC4.13	0.1458	0.0359	4.06	0.0001 ***
const	0.0017	0.0346	0.05	0.9617

Table 6.6: The first principal component's results from a second-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

The significance of the fourth component's lag three measure on the current depth bar is not entirely clear, but speaks to the temporal nature that these different vari-

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.2138	0.0179	11.96	0.0000 ***
PC2.11	0.3393	0.0208	16.31	0.0000 ***
PC3.11	-0.0999	0.0266	-3.76	0.0002 ***
PC4.11	-0.0325	0.0323	-1.01	0.3134
PC1.12	-0.0062	0.0232	-0.27	0.7875
PC2.12	-0.0465	0.0220	-2.11	0.0348 *
PC3.12	0.1137	0.0273	4.17	0.0000 ***
PC4.12	-0.0376	0.0321	-1.17	0.2423
PC1.13	-0.0133	0.0218	-0.61	0.5407
PC2.13	0.0298	0.0197	1.51	0.1308
PC3.13	-0.0001	0.0226	-0.00	0.9967
PC4.13	-0.0126	0.0303	-0.42	0.6779
const	0.0021	0.0292	0.07	0.9434

Table 6.7: The second principal component's results from a second-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

ables and consequently components have. Volume, depth, and trade count measures tend to incorporate quite a bit of information for quite an extended amount of time into its current value. The R^2 value for component one's results is relatively high at 0.28.

The shorter temporal dependence of price change measures is demonstrated by the VAR results for the second component. Component two largely captures price change measures such as the Martin Index and liquidity ratio one has high reliance on the values of components one, two, and three of lag one, and three of lag two. Furthermore, it also has slight dependence on its second-order value. This has an R^2 value of 0.18.

The estimation results for the third component indicate strong third order dependences on multiple components, specifically one, three, and four. Because component four includes composite liquidity, which is a ratio of the proportional quoted spread to the dollar depth, the dependence on the volume, trade count, and depth measure (component one) seems to rationalize this. This has an R^2 value of 0.33.

Meanwhile, the fourth component shows highly significant dependence on the first, second, and third lags of component three and itself. There is strong influence

	Estimate	Std. Error	t value	Pr(> t)	
PC1.11	0.3729	0.0141	26.38	0.0000	***
PC2.11	0.0546	0.0165	3.32	0.0009	***
PC3.11	0.2434	0.0210	11.58	0.0000	***
PC4.11	0.0240	0.0255	0.94	0.3464	
PC1.12	-0.0509	0.0183	-2.78	0.0056	**
PC2.12	0.1141	0.0174	6.55	0.0000	***
PC3.12	0.1198	0.0216	5.56	0.0000	***
PC4.12	-0.0501	0.0254	-1.97	0.0490	*
PC1.13	-0.0482	0.0173	-2.79	0.0052	**
PC2.13	0.0105	0.0156	0.67	0.5025	
PC3.13	0.0849	0.0179	4.74	0.0000	***
PC4.13	-0.1240	0.0240	-5.17	0.0000	***
const	0.0034	0.0231	0.15	0.8838	

Table 6.8: The third principal component's results from a second-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

from all of the first lag values. Because these later principal components encompass many composite measures, it seems that the earlier components, which tend to isolate categories more, influence these later components greatly. The estimation results for equation for component four have an R^2 value of 0.22.

	Estimate	Std. Error	t value	Pr(> t)	
PC1.11	0.1120	0.0114	9.86	0.0000	***
PC2.11	-0.0290	0.0132	-2.20	0.0282	*
PC3.11	-0.0979	0.0169	-5.80	0.0000	***
PC4.11	0.1630	0.0205	7.95	0.0000	***
PC1.12	-0.0122	0.0147	-0.83	0.4081	
PC2.12	-0.0502	0.0140	-3.59	0.0003	***
PC3.12	-0.0633	0.0173	-3.66	0.0003	***
PC4.12	0.1371	0.0204	6.71	0.0000	***
PC1.13	-0.0077	0.0139	-0.56	0.5769	
PC2.13	-0.0087	0.0125	-0.70	0.4867	
PC3.13	-0.0666	0.0144	-4.64	0.0000	***
PC4.13	0.0691	0.0193	3.59	0.0003	***
const	0.0011	0.0185	0.06	0.9519	

Table 6.9: The fourth principal component's results from a second-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

6.5 Financial Sector ETF Time Series Analysis

As discussed in Chapter 5.5, there are five significant principal components for the liquidity measures of the Financial Sector ETF. These components account for a total of 72.90% of the variance of the original data.

From the score plot of component one, liquidity appears to decrease and then increase over the course of the time period of the analysis, ending at roughly the same level of liquidity as where it began. Components three and five also exhibit this shape; furthermore their peaks in magnitude appear roughly around the same date, February 26, 2019. Perhaps this suggests time interdependencies between these measures. While component four's shape is slightly different, there may also be temporal interdependencies between components one, three, and four given that component four also reaches its maximum value around February 26, 2019.

Aside from the temporal nature of the scores, the magnitude of the range of the component scores decreases as the variance explained by each score decreases. This is with the exception of component five, which has the second largest range but explains the lowest variance. This is most likely because this measure trends smoothly without

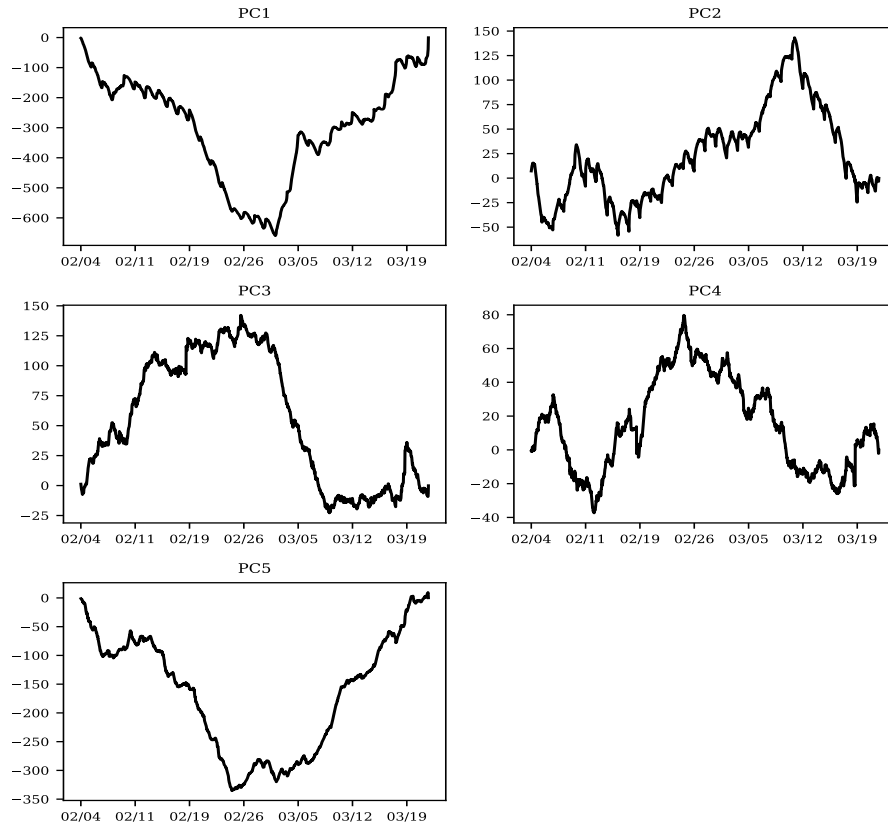


Figure 6-12: The cumulative sum of the Financial Sector ETF's significant principal components' scores over time.

much of the intraday noise that is most likely latent in other measures. Given that we have previously seen price measures to be fairly noisy, the smoothness in the steps between consecutive scores suggests it moves in somewhat consistent increments each bar.

6.5.1 Time Series Structure

From the time series of the scores, it appears that some of these components exhibit cross-correlation and most likely autocorrelation as well. Figure 6-13 displays the autocorrelations of the components and their cross-correlations.

Component one demonstrates significant autocorrelation for more than 25 lags. While all of the stocks so far have demonstrated significant autocorrelation in their volume, trade count, and depth measure, the first component here demonstrates the

strongest example of this by a significant margin. This indicates a much larger degree of momentum in the volume and depth of this ETF than a few of its underlying stocks. Beyond component one, components two and five also demonstrate autocorrelation. Again, the number of significant lags for component five is quite large, suggesting the component possesses a substantial degree of momentum. This was somewhat noticeable in the score plots, as component five's cumulative sum line was much smoother than that of the other components. The fifth component describes the contrast between liquidity ratio one, the average price change of a transaction, against the order ratio and the proportional effective spread measure. Between components, there is some significant cross-correlation, the most noticeable of which is component one lagging component two.

6.5.2 Vector Autoregressive Model

From the partial autocorrelations, a VAR(3) or even VAR(4) could be applied to components one, three, and four. However, most of the cross-correlations dependencies are only first-order. Applying the SC indicates that we can use a third-order model. The results for the first component indicate high third-order dependence on the the second component in addition to itself. Component two describes the difference between order ratio and composite liquidity. The order ratio indirectly dictates quote sizes and the size of future trades, so this dependence seems fitting. Because volume, trade count measures, and depth have shown to have a strong dependence on their previous values, the results for this component are in line with those of Bank of America and Citigroup. The R^2 value is quite high at 0.38.

The second component has similar dependencies to the first component with the exception that it is influenced significantly by all components from the previous bar. Because this second component largely describes a a value that is ultimately contrived by market makers, it also seems to be systematic, that whenever measures are analyzed that are completely reliant on market makers, they use all of the information available to them to adjust their internal models for quotes. The R^2 value is quite similar to that of component one at 0.40, which indicates a 63% chance of correct

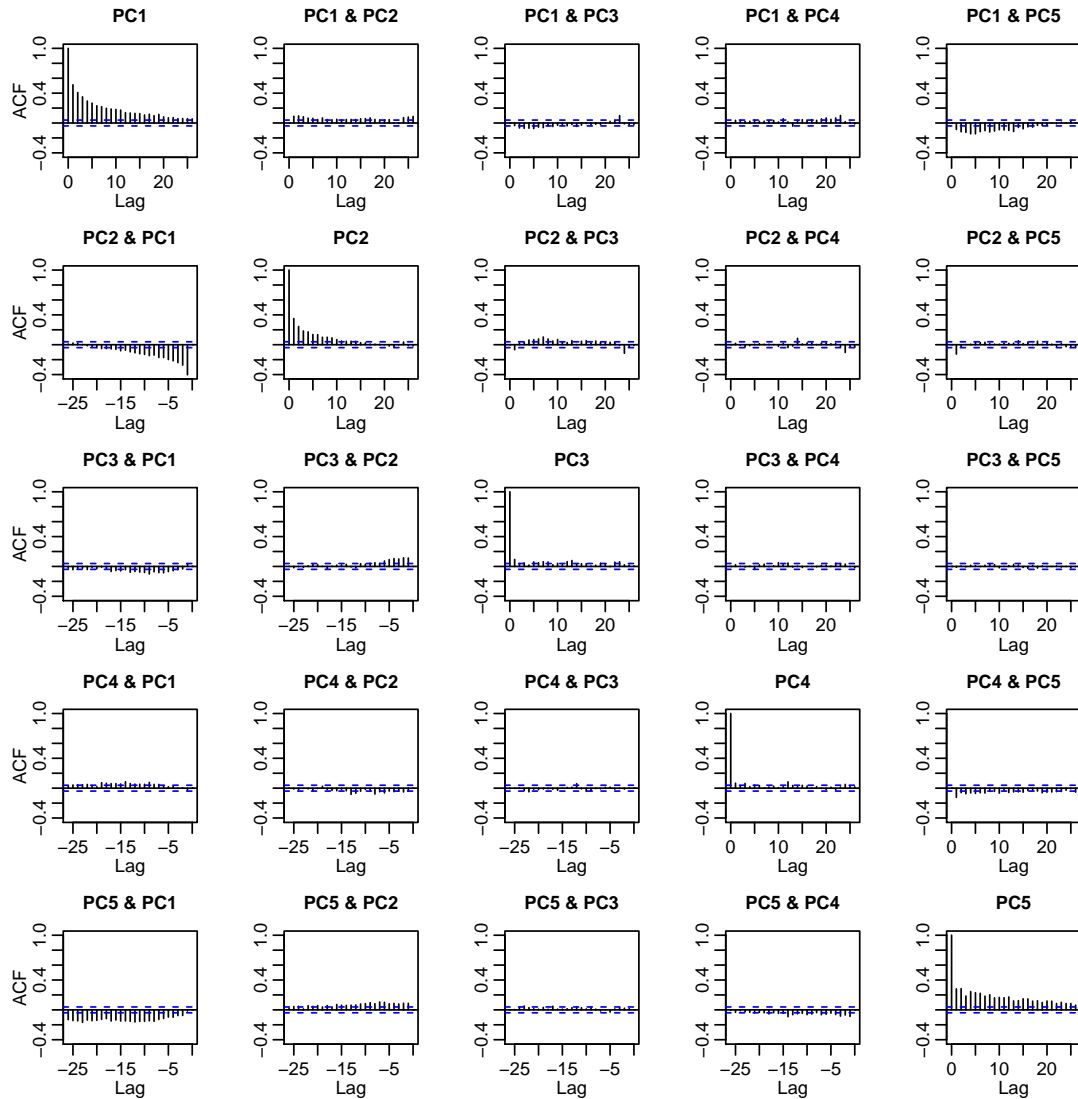


Figure 6-13: The autocorrelation and cross-correlation of the Financial Sector ETF's principal components scores over time.

forecasting.

The results for the third and fourth component are shown in Tables 6.12 and 6.13. Their dependencies in general are much weaker than those for components one and two, as indicated both earlier when examining their lack of autocorrelation and time series structure, and as well as considering the R^2 values of the estimations for these components, which are 0.04 and 0.03 respectively.

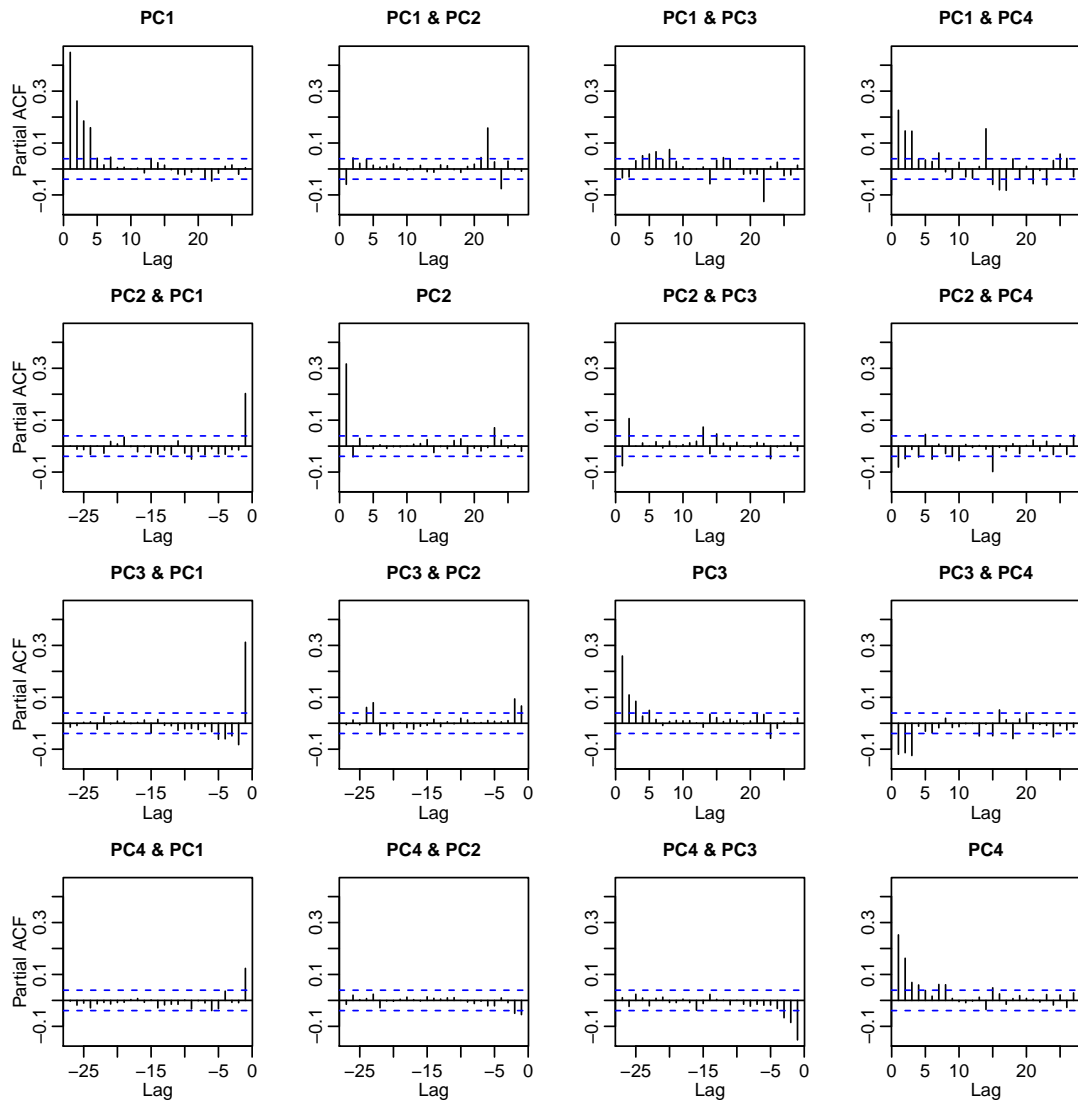


Figure 6-14: The partial autocorrelation and partial cross-correlation of the Financial Sector ETF’s principal components scores over time.

Component five shows more time series structure than components three and four, but its structure is still not as significant as those of components one and two. As postulated earlier, component five seems to have a strong underlying momentum to it. It is highly statistically influenced by its previous three values. The coefficients are such that the in addition to the other components, the score for component five equals roughly the average of the last two of its value with 0.07 of its measure from

	Estimate	Std. Error	t value	Pr(> t)	
PC1.11	0.2949	0.0209	14.12	0.0000	***
PC2.11	0.4179	0.0366	11.43	0.0000	***
PC3.11	-0.0549	0.0302	-1.82	0.0691	
PC4.11	0.0706	0.0315	2.24	0.0252	*
PC5.11	-0.0976	0.0357	-2.74	0.0062	**
PC1.12	0.2817	0.0245	11.48	0.0000	***
PC2.12	0.1567	0.0381	4.12	0.0000	***
PC3.12	-0.0461	0.0305	-1.51	0.1307	
PC4.12	0.0285	0.0317	0.90	0.3682	
PC5.12	-0.0337	0.0363	-0.93	0.3540	
PC1.13	0.2113	0.0236	8.96	0.0000	***
PC2.13	-0.0842	0.0317	-2.65	0.0080	**
PC3.13	-0.0504	0.0301	-1.68	0.0940	
PC4.13	-0.0139	0.0314	-0.44	0.6592	
PC5.13	-0.0365	0.0362	-1.01	0.3130	
const	-0.0083	0.0335	-0.25	0.8049	

Table 6.10: The first principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

fifteen minutes earlier. Its R^2 value is 0.16.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	-0.3071	0.0116	-26.38	0.0000 ***
PC2.11	0.3008	0.0204	14.76	0.0000 ***
PC3.11	-0.1343	0.0168	-7.98	0.0000 ***
PC4.11	0.0479	0.0176	2.72	0.0065 **
PC5.11	-0.2116	0.0199	-10.65	0.0000 ***
PC1.12	0.0269	0.0137	1.97	0.0490 *
PC2.12	0.2423	0.0212	11.42	0.0000 ***
PC3.12	0.0239	0.0170	1.41	0.1592
PC4.12	0.0084	0.0177	0.48	0.6335
PC5.12	-0.0102	0.0202	-0.50	0.6141
PC1.13	0.0667	0.0131	5.08	0.0000 ***
PC2.13	0.1005	0.0177	5.69	0.0000 ***
PC3.13	0.0197	0.0168	1.17	0.2409
PC4.13	-0.0384	0.0175	-2.19	0.0285 *
PC5.13	0.0541	0.0202	2.68	0.0074 **
const	0.0004	0.0187	0.02	0.9826

Table 6.11: The second principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.0331	0.0140	2.36	0.0181 *
PC2.11	0.0557	0.0245	2.28	0.0229 *
PC3.11	0.0786	0.0202	3.89	0.0001 ***
PC4.11	0.0339	0.0211	1.60	0.1088
PC5.11	-0.0362	0.0239	-1.52	0.1290
PC1.12	-0.0248	0.0164	-1.51	0.1316
PC2.12	0.0643	0.0255	2.53	0.0116 *
PC3.12	0.0277	0.0204	1.36	0.1750
PC4.12	-0.0021	0.0212	-0.10	0.9203
PC5.12	0.0283	0.0243	1.16	0.2442
PC1.13	-0.0002	0.0158	-0.02	0.9874
PC2.13	0.0471	0.0212	2.22	0.0264 *
PC3.13	0.0483	0.0201	2.40	0.0164 *
PC4.13	-0.0108	0.0210	-0.51	0.6072
PC5.13	0.0123	0.0242	0.51	0.6120
const	-0.0007	0.0225	-0.03	0.9755

Table 6.12: The third principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	-0.0162	0.0133	-1.21	0.2247
PC2.11	-0.0282	0.0233	-1.21	0.2272
PC3.11	0.0010	0.0192	0.05	0.9600
PC4.11	0.0592	0.0201	2.95	0.0032 **
PC5.11	-0.1056	0.0227	-4.65	0.0000 ***
PC1.12	0.0090	0.0156	0.58	0.5651
PC2.12	-0.0432	0.0243	-1.78	0.0753
PC3.12	-0.0143	0.0194	-0.73	0.4634
PC4.12	0.0055	0.0202	0.27	0.7844
PC5.12	-0.0148	0.0231	-0.64	0.5220
PC1.13	-0.0195	0.0150	-1.29	0.1959
PC2.13	0.0003	0.0202	0.01	0.9886
PC3.13	-0.0046	0.0192	-0.24	0.8093
PC4.13	0.0547	0.0200	2.73	0.0063 **
PC5.13	-0.0523	0.0231	-2.27	0.0236 *
const	0.0012	0.0214	0.06	0.9540

Table 6.13: The fourth principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.0242	0.0119	2.04	0.0413 *
PC2.11	0.0756	0.0207	3.65	0.0003 ***
PC3.11	0.0165	0.0171	0.96	0.3369
PC4.11	-0.0507	0.0179	-2.84	0.0046 **
PC5.11	0.1906	0.0202	9.42	0.0000 ***
PC1.12	-0.0173	0.0139	-1.24	0.2136
PC2.12	0.0322	0.0216	1.49	0.1365
PC3.12	0.0270	0.0173	1.56	0.1184
PC4.12	-0.0399	0.0180	-2.22	0.0264 *
PC5.12	0.2198	0.0206	10.67	0.0000 ***
PC1.13	-0.0071	0.0134	-0.53	0.5963
PC2.13	-0.0034	0.0180	-0.19	0.8480
PC3.13	0.0214	0.0171	1.26	0.2090
PC4.13	-0.0460	0.0178	-2.58	0.0099**
PC5.13	0.0739	0.0205	3.60	0.0003***
const	-0.0008	0.0190	-0.04	0.9667

Table 6.14: The fifth principal component's results from a third-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

Chapter 7

Principal Components Analysis of Daily-normalized Liquidity Measures

As demonstrated in Chapter 4, a collection of the liquidity measures demonstrate clear intraday liquidity patterns for all securities analyzed. In order to investigate how these measures behave on an intraday basis, rather than the long-term analysis captured in Chapter 5, we perform a second principal components analysis on the same thirteen liquidity measures from Chapter 5, with the exception that these measures are now normalized daily rather than over the entire time period of 32 days. Daily normalization enforces second-order stationarity of the liquidity measures because each day's mean level is zero and each day's standard deviation is one. This normalization enables the analysis to focus on models of the intraday dynamics of the measures.

7.1 Principal Components Analysis of Bank of America's Daily-normalized Liquidity Measures

The liquidity measure data for Bank of America consists of 2,483 observations of 13 liquidity measures. These 13 measures are standardized by day so that each measure has a mean of zero and standard deviation of one each day. Table 7.1 contains the significant principal components and Figure 7-1 shows the scree plot of the eigenvalues from two sets of principal component analyses: one where the liquidity measures are normalized daily and the other where the measures are normalized over the 32-day analysis period.

In terms of choosing significant components, the Kaiser criterion indicates that the first four components are considered meaningful. However, because the fifth component has an eigenvalue very close to one and accounts for 71.45% of the variance of the proportional realized spread, a variable whose variance is not sufficiently explained by the first four components, we choose to include the fifth component in our analysis.

Overall, the amount of variance explained by the principal components of the daily measures is incredibly similar to those found in Chapter 5.3. The first component accounts for 35.95% of the variance, which is quite large. Components two and three cover similar amounts of variances at 13.65% and 11.51% respectively. Finally the last two components also explain similar variances of around 8%.

	Standard deviation	Variance	Proportion of Variance	Cumulative Proportion
PC1	2.161697	4.672935	35.95%	35.95%
PC2	1.331980	1.774171	13.65%	49.59%
PC3	1.223437	1.496799	11.51%	61.11%
PC4	1.036741	1.074832	8.27%	69.37%
PC5	0.989595	0.979298	7.53%	76.91%

Table 7.1: The first five principal components of a principal components analysis on Bank of America's 13 daily-normalized liquidity measures.




Figure 7-1: The scree plot of Bank of America’s principal components resulting from two principal component analyses; one on daily-normalized measures and the other on measures normalized over a 32-day time period.

The similarity of the amounts of variances explained by each principal component suggests that there may be parallels between the components of the 32-day normalized measures and the components of the daily-normalized measures. The first principal component demonstrates this as it is effectively the same as the first component from the analysis in Chapter 5.3. More specifically, component one in this analysis is a weighted average of log depth, dollar depth, trade count, turnover, and flow ratio, which is the ratio of turnover to average time in between trades. These results and those for the remaining significant components are shown in Figure 7-2 and Table 7.2.

The second component is largely a weighted average of the proportional effective spread, composite liquidity, and the Martin Index. The proportional effective spread




Figure 7-2: Bank of America's principal component loadings from principal components analysis on 13 daily-normalized liquidity measures.

captures the spread from mid at which a trade occurs, while the Martin Index divides the squared price change between consecutive trades by turnover. Composite liquidity is the proportional quoted spread over dollar depth. As described earlier when observing a positive correlation between the Martin Index and the proportional effective spread, the weighted sum of the measures describes the variance of price changes between trades and the spread from mid at which a trade occurs. The positive correlation between these measures may initially seem odd; however, when the effective spread, (the spread from the mid-price at which trades occur) is wider (tighter), the price changes between the current trades and previous trades are greater (less). In a way, the value of the Martin Index is in part due to the behavior of the proportional effective spread.

Component three is a weighted average of the proportional quoted spread and the order ratio, which is the imbalance in quoted bid and ask sizes divided by turnover. Thus this component can be thought of as a weighted average of the quoted depth imbalance and the quoted spread. To simplify further, this component might be conceptualized as an indicator of the behavior of market makers as it encompasses the two tools they can use to skew markets and influence investor behavior, depth

	PC1	PC2	PC3	PC4	PC5
PQS	-0.001	-0.028	-0.450	-0.631	-0.289
PES	-0.023	-0.512	-0.302	0.214	0.034
PRS	-0.008	0.125	-0.114	0.316	-0.854
T	0.435	0.171	-0.056	0.048	0.028
\$D	0.398	-0.193	0.244	0.034	-0.045
LogD	0.395	-0.245	0.126	-0.138	-0.122
CL	-0.080	-0.559	0.272	-0.289	-0.203
OR	0.014	0.223	-0.537	-0.055	0.174
N	0.446	0.110	-0.031	0.022	0.005
FR	0.423	0.125	0.063	0.130	0.044
M	0.008	-0.406	-0.367	0.504	0.133
LR1	0.309	-0.125	-0.327	-0.107	-0.082
IlliQ	0.122	-0.157	-0.081	-0.254	0.265

Table 7.2: The principal component loadings of each of Bank of America’s 13 daily-normalized liquidity measures across components.

imbalance and quoted spread.

The fourth component captures the contrast between the proportional quoted spread and the Martin Index. The negative correlation between the two measures indicates that as the price changes between transactions increase or the turnover in the bar decreases significantly compared to the price changes, indicating decreasing liquidity, then the quoted spread decreases, indicating rising liquidity. Perhaps the unexpected behavior of the quoted spread in this scenario is a result of the tick size restricting Bank of America’s spread from accurately reflecting changes in liquidity conditions.

Finally, the fifth component is in a large part determined by the value of the proportional realized spread, as its loading, -0.85, dwarfs the other liquidity measures’ loadings, the next largest of which are the quoted spread at -0.29 and ILLIQ at 0.265.

In the varimax rotation in Table 7.3, the meanings of these components change slightly. While components one, two, and five remain the same, the meanings of the third and fourth components shift slightly. the third component evolves into a contrast of composite liquidity and log depth against the order ratio. This interpretation of the third component as a contrast between composite liquidity and order ratio matches the variables that describe the third component in the first principal com-

ponents analysis. In other words, the first three components in this varimax rotation have the same themes and variables as they did in the principal components analysis of the liquidity measures standardized over the analysis time period.

The varimax rotation turns The fourth component into an average of the proportional quoted spread, liquidity ratio one, an order ratio. This involves the quoted spread and depth in relation with the average cost per transaction. Finally, the fifth component contrasts proportional realized spread against ILLIQ, or more tractably, the change in price across five minute bars compared to that change adjusted for the volume traded during that period.

	PC1	PC2	PC3	PC4	PC5
PQS				-0.819	
PES		-0.612			
PRS					-0.925
LogD	0.412		0.298		
\$D	0.42				
N	0.439				
T	0.424				
CL			0.705		
M		-0.738			
LR1	0.294			-0.325	
FR	0.421				
OR			-0.487	-0.318	
Illiq					0.356

Table 7.3: A varimax rotation of Bank of America’s first five principal component loadings resulting from principal components analysis on daily-normalized liquidity measures.

Table 7.4 displays the original variable measures and the percentage of their variance that is accounted for by each principal component. The components best explain volume, depth, and trade count measures, with trade count being the most well explained variable at 95.34% of its variance explained. On the other end of the spectrum are ILLIQ and order ratio. ILLIQ is split somewhat equally across components one, three, four, and five showing that its essence is not truly being captured by any one of these components. However, that is not the case for order ratio, as 43.18% of its variance is described by component three.

	PC1	PC2	PC3	PC4	PC5	Cumulative Variance Explained
N	92.99%	2.15%	0.14%	0.05%	0.00%	95.34%
T	88.46%	5.19%	0.47%	0.25%	0.08%	94.44%
\$D	74.05%	6.61%	8.91%	0.12%	0.20%	89.90%
LogD	72.94%	10.65%	2.38%	2.05%	1.46%	89.48%
FR	83.65%	2.77%	0.59%	1.82%	0.19%	89.02%
PRS	0.03%	2.77%	1.95%	10.74%	71.45%	86.94%
CL	2.99%	55.46%	11.08%	8.98%	4.04%	82.55%
PQS	0.00%	0.14%	30.32%	42.81%	8.18%	81.46%
M	0.03%	29.26%	20.17%	27.31%	1.73%	78.50%
PES	0.25%	46.53%	13.66%	4.92%	0.11%	65.47%
LR1	44.64%	2.77%	16.01%	1.23%	0.66%	65.31%
OR	0.09%	8.83%	43.18%	0.33%	2.97%	55.39%
Illiq	6.96%	4.37%	0.98%	6.94%	6.88%	26.13%

Table 7.4: Percentage of cumulative variance of each of BAC’s original liquidity measures as explained by the five significant PCA components.

7.2 Principal Component Analysis of Citigroup’s Daily-normalized Liquidity Measures

Extending upon the initial overall liquidity measure analysis in Chapter 5.4, we apply principal component analysis to the daily normalized five-minute bar aggregated measures. As with the original analysis, there are 13 measures and 2,485 five-minute bar observations, split across 32 trading days. The Kaiser criterion identifies the four principal components as significant; explaining a total of 76.27% of the original variance in the data as indicated in Figure 7.5.

	Standard deviation	Variance	Proportion of Variance	Cumulative Proportion
PC1	2.153910	4.639330	35.69%	35.69%
PC2	1.637352	2.680923	20.62%	56.31%
PC3	1.245482	1.551226	11.93%	68.24%
PC4	1.021876	1.044231	8.03%	76.27%

Table 7.5: The principal components of Citigroup using PCA on 13 daily-normalized liquidity measures and the Kaiser criterion to determine significance.

When comparing the amounts of variances explained by each of these principal components to those on the overall liquidity measures, the variances are very similar;

these similarities can be seen in the scree plot that contains the components for both principal component analyses. On the daily measures, component one explains more variance, while on the overall PCA, component three accounts for more variance than component three on the daily measures.




Figure 7-3: The scree plot of Citigroup’s principal components resulting from two principal component analyses on 13 liquidity measures normalized daily and normalized over the entire 32-day time period.

The loadings of principal component analysis on the daily liquidity measures is displayed in Table 7-4. As is familiar, the first component is an average of the volume, depth, and trade count measures log depth, dollar depth, trade count, turnover, and flow ratio. This is effectively the same component as the analysis of overall liquidity measures produces; however, it accounts for about 4% more variance of the data. Perhaps this suggests that these measures are more easily able to be modeled on an intraday basis given their systematic intraday liquidity pattern, as shown in Chapter

4.2.




Figure 7-4: The principal component loadings of Citigroup resulting from PCA on 13 liquidity measures.

The second principal component explains 20.62% of the variance, and is a weighted average of the proportional quoted spread, the proportional effective spread, and composite liquidity. This measure captures the spread between the best bid and ask that is being shown, as well as this spread over the dollar depth, and finally the price from at which a trade occurs. It is intuitive and expected that these measures are correlated.

The third component is a weighted average of the Martin Index, liquidity ratio one, and order ratio with the Martin Index being weighted the most at -0.515, liquidity ratio one weighted the second most at -0.479, and finally the order ratio with a loading of -0.43. Thus this component describes price change measures between consecutive trades over trade count and turnover with some element of the imbalance between the quoted bid size and the quoted ask size. Because when liquidity decreases, price measures change, and given that Citigroup's spread does not often adjust for different levels of liquidity (as shown in Chapter 4.2), the adjustment often comes through a change in bid and ask sizes. Perhaps this explains why these particular measures are grouped together.

The fourth component describes the contrast between the proportional realized spread and the order ratio, and in essence, captures the difference between realized five-minute market impact and the proportional quoted spread over dollar depth. When the price changes more over a five-minute bar, the order ratio tends to reflect higher liquidity, even though larger price changes over a five-minute period suggest the opposite. A possible explanation for this behavior could be that the absolute quoted sizes both decrease and therefore their difference does as well, producing an order ratio that indicates higher liquidity. Alternatively, as the price changes over five-minutes grow smaller, quoted sizes increase and therefore so does their absolute spread.

	PC1	PC2	PC3	PC4
PQS	-0.113	-0.453	0.178	-0.323
PES	0.020	-0.507	-0.266	0.034
PRS	0.028	0.042	0.033	0.731
T	0.431	0.157	-0.069	-0.073
\$D	0.421	-0.098	0.214	-0.023
LogD	0.394	-0.178	0.173	0.021
CL	-0.067	-0.504	0.360	-0.116
OR	-0.089	0.189	-0.430	-0.431
N	0.445	0.066	-0.034	-0.136
FR	0.434	0.075	0.061	-0.090
M	0.026	-0.314	-0.515	0.221
LR1	0.248	-0.228	-0.479	-0.006
Illiq	0.086	-0.144	0.025	0.280

Table 7.6: The principal component loadings of each of Citigroup’s 13 liquidity measures across components.

The varimax rotation of these loadings, displayed in Table 7.6 presents a similar albeit slightly different interpretation of these components. While component one still remains the same, component two is no longer influenced by the effective spread and is solely a weighted average of the proportional quoted spread and composite liquidity. This simplifies the component but does not necessarily change the meaning behind it.

The varimax rotation moves the assignment of the proportional effective spread to component three instead. After the varimax rotation, component three now reflects a

	PC1	PC2	PC3	PC4
PQS		-0.547		
PES			-0.506	
PRS				0.703
LogD	0.411			
\$D	0.45			
N	0.447			
T	0.42			
CL		-0.632		
M			-0.627	
LR1			-0.546	
FR	0.446			
OR				-0.556
Illiq				0.293

Table 7.7: A varimax rotation of Citigroup’s first four principal component loadings resulting from principal components analysis on daily-normalized liquidity measures.

weighted average of the Martin Index, liquidity ratio one, and the proportional effective spread. Because the Martin Index and liquidity ratio one both use consecutive trade prices in its calculation, they are naturally related to consecutive measurements of the proportional effective spread, which use the trade price at time t in its calculation.

Finally, after the varimax rotation, component five represents the order ratio subtracted from a weighted sum of the proportional realized spread and ILLIQ. The proportional realized spread and ILLIQ, the log return of the five-minute bar divided by turnover, are both measures of five minute prices changes. The way that the order ratio responds to price changes, specifically increases or decreases in liquidity is one that seems to arise frequently. While price changes may increase, quoted depth often correspondingly decreases, causing the absolute difference between quoted bid and ask sizes to decrease and the order ratio to decrease, artificially indicating better liquidity.

As far as which variables are best explained on an intraday basis, the percentages of the original variables explained by the components are contained in Table 7.8. The depth and trade count measures are the most well explained while a mixture of spread and price change measures are the least explained. ILLIQ is particularly

unexplained, where cumulatively only 17.28% of the variance is captured across the first five components.

When comparing these to the percentages of the overall measures explained by the first five components, perhaps we can gain insight into the natural time horizons of these measures. Presumably, longer-term trend liquidity measures might be better explained by the PCA of the overall measures, while measures with strong intraday measures may be more suited for the PCA on the daily liquidity measures. Table 5.12 includes the corresponding percentages of each overall liquidity measure explained by the significant components from the analysis in Chapter 5.

	PC1	PC2	PC3	PC4	Cumulative Variance Explained
N	91.91%	1.17%	0.18%	1.93%	95.19%
T	86.22%	6.61%	0.74%	0.56%	94.12%
\$D	82.26%	2.58%	7.11%	0.06%	92.00%
CL	2.08%	68.13%	20.11%	1.41%	91.73%
FR	87.42%	1.51%	0.58%	0.85%	90.35%
LogD	72.05%	8.50%	4.64%	0.05%	85.24%
PES	0.19%	68.94%	10.98%	0.12%	80.23%
LR1	28.55%	13.94%	35.61%	0.00%	78.10%
PQS	5.93%	55.04%	4.92%	10.90%	76.78%
M	0.31%	26.44%	41.16%	5.10%	73.02%
OR	3.68%	9.58%	28.69%	19.41%	61.36%
PRS	0.36%	0.47%	0.17%	55.82%	56.83%
Illiq	3.43%	5.56%	0.10%	8.19%	17.28%

Table 7.8: Percentage of cumulative variance of each of Citigroup’s daily liquidity measures as explained by the four significant PCA components.

The largest similarity between the two is that trade count, volume, and depth measures are well explained by both analyses, with the variance of these types of measures slightly more accounted for in the daily analysis. ILLIQ becomes dramatically less explained in the daily analysis, as it goes from 59% in the overall analysis to just 17% here. ILLIQ is the one measure used whose intended use is a longer term, and the empirical evidence appears to support this. Composite liquidity is better explained on a daily basis, rather than long term. Otherwise, the measures have relatively comparable variances explained by the different analyses.

7.3 Principal Components Analysis of the Financial Sector ETF's Daily-normalized Liquidity Measures

This principal components analysis of the XLF ETF is performed on 2,496 daily-normalized observations of 13 liquidity measures. Table 7.9 shows the components we consider significant for this analysis. Although the the Kaiser criterion designates only the first four components as significant (since these are the only components with eigenvalues greater than one), because four components only account for 66.99% of the variance of the data and because with an eigenvalue of 0.9999, component five is incredibly close to being considered significant by the Kaiser criterion, we choose to consider the first components as significant in this analysis.

	Standard deviation	Variance	Proportion of Variance	Cumulative Proportion
PC1	2.171561	4.715677	36.27%	36.27%
PC2	1.308260	1.711543	13.17%	49.44%
PC3	1.105219	1.221509	9.40%	58.84%
PC4	1.029522	1.059915	8.15%	66.99%
PC5	0.999957	0.999914	7.69%	74.68%

Table 7.9: The principal components of the Financial Sector ETF resulting from principal components analysis on 2,496 observations of 13 daily-normalized liquidity measures.

The amounts of variances explained by the components for the individual liquidity measures are extremely close to those explained by corresponding components in the earlier analysis of the XLF ETF's liquidity measures. A scree plot of the eigenvalues from each analysis is shown in Figure 7-5. Across components, the variances are effectively equal, with perhaps the most notable differences in components two and four. Component two captures more variance in the analysis of the daily measures, while component four of the 32-day normalized data describes more variance.

Table 7.10 and Figure 7-6 shows the loadings of the principal components of the Financial Sector ETF's daily-normalized liquidity measures. The first component is an average of log depth, dollar depth, trade count, turnover, and flow ratio, and it




Figure 7-5: The scree plots of the Financial Sector ETF’s principal components resulting from two principal component analyses, one on 2,946 observations of 13 liquidity measures each, where one analysis normalizes the measures over the entire time period (32 days) and the other normalizes them daily.

accounts for 36.27% of the variance. This first components’ loadings and variance accounted for are effectively equal to those suggests that perhaps not much further insights is gained from changing the context in which these volume, depth, and trade count measures are examined.

The second principal component is a weighted average of composite liquidity and the Martin Index contrasted against the quoted spread, which suggests the component encompasses the price change between consecutive trades and how the dynamics between quoted spread and depth respond to this compared to just how the quoted spread responds. If we consider the Martin Index to be representative of large (or small) price movements given a certain amount of volume, then in response to this




Figure 7-6: The principal component loadings of the Financial Sector ETF resulting from PCA on 2,496 daily-normalized observations of 13 liquidity measures.

change in liquidity, as indicated by the changes in prices, then we might expect a market maker to respond to this by changing the quoted spread and/or the quoted depth.

If these two levers (the quoted spread and quoted depth) do not move in conjunction with each other, then potentially this could be responsible for the negative correlation between composite liquidity and the quoted spread. The Martin Index and the composite liquidity being positive correlated indicates that when price changes increase, suggesting worsening liquidity, the composite liquidity decreases while the quoted spread decreases or remains the same. This indicates that in response to changing liquidity conditions, the size of the quoted depth changing may be more reliable. The positive correlation between dollar depth, log depth, and composite liquidity support this notion that the measure is moving more in conjunction with depth rather than spread.

The third component is a weighted average of the proportional quoted spread, the proportional effective spread, and the Martin Index, effectively making it a price change measure that is partly weighted by spread measures. The fourth component describes the contrast between two sets, each of which include two measures. In

	PC1	PC2	PC3	PC4	PC5
PQS	-0.026	0.308	-0.584	-0.057	0.298
PES	0.141	-0.284	-0.437	0.479	-0.099
PRS	0.009	0.184	-0.033	0.454	0.445
T	0.409	0.192	0.009	-0.015	-0.043
\$D	0.432	-0.125	0.125	0.150	-0.009
LogD	0.392	-0.254	0.054	-0.101	0.073
CL	-0.005	-0.618	-0.019	0.049	-0.041
OR	0.047	0.289	-0.404	0.171	-0.590
N	0.441	0.123	0.071	0.053	0.037
FR	0.418	0.094	0.108	0.202	-0.046
M	0.008	-0.424	-0.442	-0.101	0.156
LR1	0.279	0.066	-0.228	-0.501	0.375
Illiq	0.153	-0.026	-0.144	-0.439	-0.421

Table 7.10: The principal component loadings of each of the Financial Sector ETF's 13 liquidity measures across components.

this component, the proportional effective and the proportional realized spreads are contrasted against liquidity ratio one and ILLIQ. In this sense, together liquidity ratio one and ILLIQ encompass almost all of the various dimensions of liquidity related to price changes, as liquidity ratio one includes price changes between consecutive trades divided by trade count and ILLIQ measures the longer-term five-minute bar log return over turnover. Perhaps the contrast between these price change measures and the spread measures are an indicator of the degree to which the spread is responding to these liquidity changes.

The fifth component covers the difference between liquidity ratio one and the proportional realized spread against ILLIQ. This contrasts the average price change of a transaction and the absolute price change over a five minute bar to the log return over a five-minute bar divided by turnover. Because price changes can be scaled in different ways, this fifth component seems to compare them.

The varimax rotation of these loadings is shown in Table 7.11, and while in some components, a few variables assigned with more weight to certain components, for the most part, the intention of the measure remains the same as in the original, unrotated loadings matrix.

The amount of variance of each measure explained by the significant components

	PC1	PC2	PC3	PC4	PC5
PQS			-0.646		
PES		-0.601			-0.328
PRS				0.649	
LogD	0.362				
\$D	0.459				
N	0.454				
T	0.405				
CL		-0.517			
M		-0.544	-0.284		
LR1			-0.629		
FR	0.461				
OR					-0.781
IlliQ				-0.601	

Table 7.11: A varimax rotation of the Financial Sector ETF's first five principal component loadings resulting from principal components analysis on daily-normalized liquidity measures.

is shown in Table 7.12. The lowest ranked variables appear to be those that have moderate but not necessarily significant loadings across many component. This could be generally said for the proportional quoted spread, the proportional effective spread, the proportional realized spread, the Martin Index, ILLIQ, composite liquidity, and liquidity ratio one, all of which appear to be the variables that are least explained by the first five principal components.

	PC1	PC2	PC3	PC4	PC5	Cumulative Variance Explained
N	91.75%	2.59%	0.62%	0.30%	0.14%	95.39%
\$D	88.04%	2.68%	1.91%	2.39%	0.01%	95.02%
FR	82.43%	1.51%	1.43%	4.33%	0.21%	89.90%
LogD	72.49%	11.05%	0.36%	1.08%	0.53%	85.51%
T	78.92%	6.31%	0.01%	0.02%	0.18%	85.45%
LR1	36.72%	0.75%	6.35%	26.61%	14.07%	84.50%
OR	1.04%	14.30%	19.94%	3.10%	34.82%	73.21%
PES	9.38%	13.81%	23.34%	24.33%	0.98%	71.83%
PQS	0.32%	16.24%	41.68%	0.34%	8.88%	67.47%
CL	0.01%	65.39%	0.04%	0.25%	0.17%	65.87%
M	0.03%	30.78%	23.87%	1.08%	2.43%	58.20%
Illiq	11.04%	0.12%	2.53%	20.43%	17.73%	51.86%
PRS	0.04%	5.80%	0.13%	21.86%	19.81%	47.63%

Table 7.12: Percentage of cumulative variance of each of the Financial Sector ETF's daily-normalized liquidity measures as explained by principal components analysis.

Chapter 8

Time Series Analysis of Principal Component Scores of Daily-normalized Liquidity Measures

Because principal component analysis is an analysis of spatial variance, time is not explicitly captured in the analysis of the daily-normalized liquidity measures in Chapter 7. Even so, the results from the analysis still suggest there is a temporal nature to the component scores. While these temporal patterns were analyzed in Chapter 6, the analysis was done in the context of the 32-day time period. Applying VAR models to the principal component scores of the daily-normalized liquidity measures will give insight into the temporal relationships of these scores on an intraday basis.

8.1 Bank of America Time Series Analysis

Although there were a few differences between the principal component loadings and results between the analysis done on the liquidity measures normalized over a 32-day period compared to those normalized daily, generally, the results were fairly

consistent.

However, when examining the cumulative sums of Bank of America's principal component scores of its daily-normalized liquidity measures plotted over the time in Figure 8-1; the differences between the two analyses becomes apparent. Figure 8-1 clearly shows time-of-day dependencies of these components. The seasonality is evident in every score, even components that explain less variance such as components three and four. Component one, the depth measure, exhibits the U-shaped pattern every day throughout the time series. This suggests the days where the U-shape is smaller than others in magnitude or lower in absolute value such as February 20, 2019 have lower liquidity. Component two contains an average of the price change between consecutive trades squared divided by volume and the spread from mid that a trade occurs at divided by the mid price. Essentially capturing the price change between trades and the price in general, the same U-shape is present each day. This again suggests low liquidity in the middle of the day (since the loadings of the dominant variables are negative). Component three, an average measure of spread and quoted depth, also appears in a daily U-shape pattern, which fits with research demonstrating this effect. The fourth and fifth components still have the vague sense of a U-shape but less clearly. Perhaps this is because they capture more composite measures and factors such as longer-term five-minute price movement.

8.1.1 Time Series Structure

The seasonality of these scores is further evident through its autocorrelations and partial autocorrelations in Figure 8-2 and Figure 8-3. Component one, the depth, volume, and trade count component, demonstrates significant autocorrelation with significant lags until 15. Components two and three demonstrate some autocorrelation but it is rather weak and is only for two to three lags at most. Components four and five do not show any significant indications of autocorrelation.

Meanwhile, there is significant cross-correlation between component one and every other significant component. These cross-correlations are such that components two through five have values that help predict more immediate values of component one,

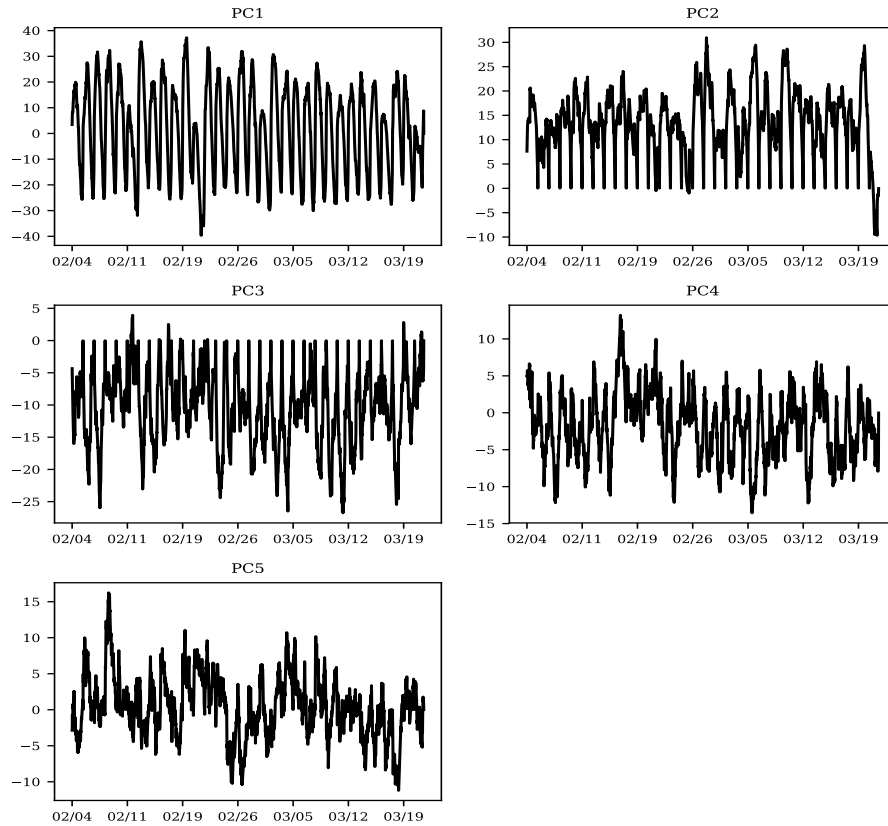


Figure 8-1: The cumulative sum of BAC's principal components scores over time.

turnover, depth, and trade count. However, in turn, component one's current values influence components two through five's values on a longer time scale into the future. From the most significant lag correlations, component two's current scores appear to have the most influential effect on how volume, depth, and trade count will act in the short-term over the next five to thirty minutes. However, on the longer time horizon, such as an hour, component four's values influence component one's current scores more.

The partial autocorrelations in Figure 8-3 illustrate that most of these components are VAR(1) with the exception of component one which could be as high as VAR(5).

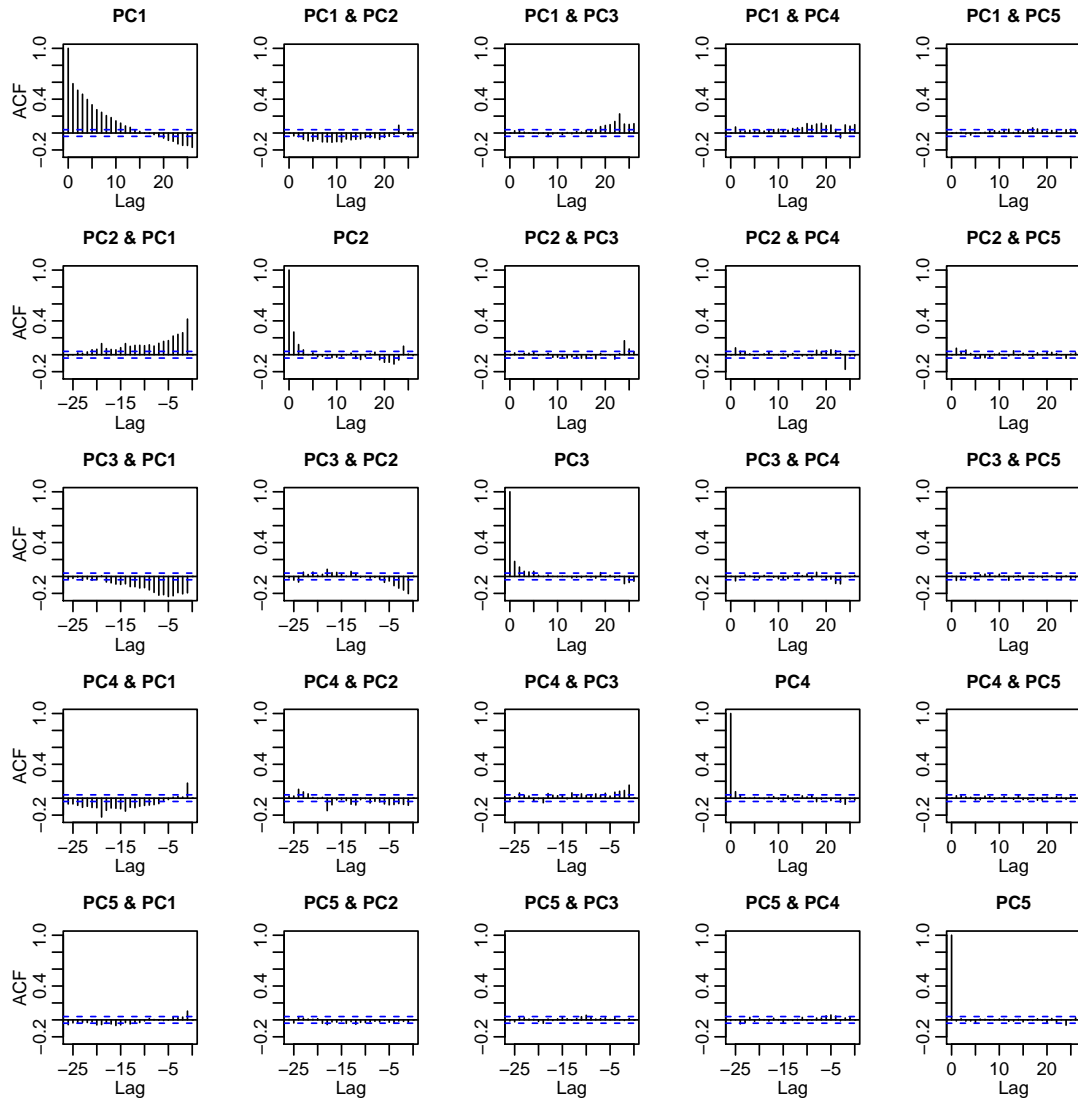


Figure 8-2: The autocorrelation and cross-correlation of Bank of America’s principal components scores of Daily-normalized Liquidity measures.

8.1.2 Vector Autoregressive Model

Applying the Schwarz Criterion (SC) to these scores results in a second-order VAR model. Fitting the model to these scores produces the following coefficient results.

The leads and lags as discussed above are further confirmed by the significance of some components’ lags. Component one’s R^2 is very high at 0.43, indicating an R value of 0.65. Component two’s R^2 is 0.31. Component three’s R^2 is smaller at 0.14,

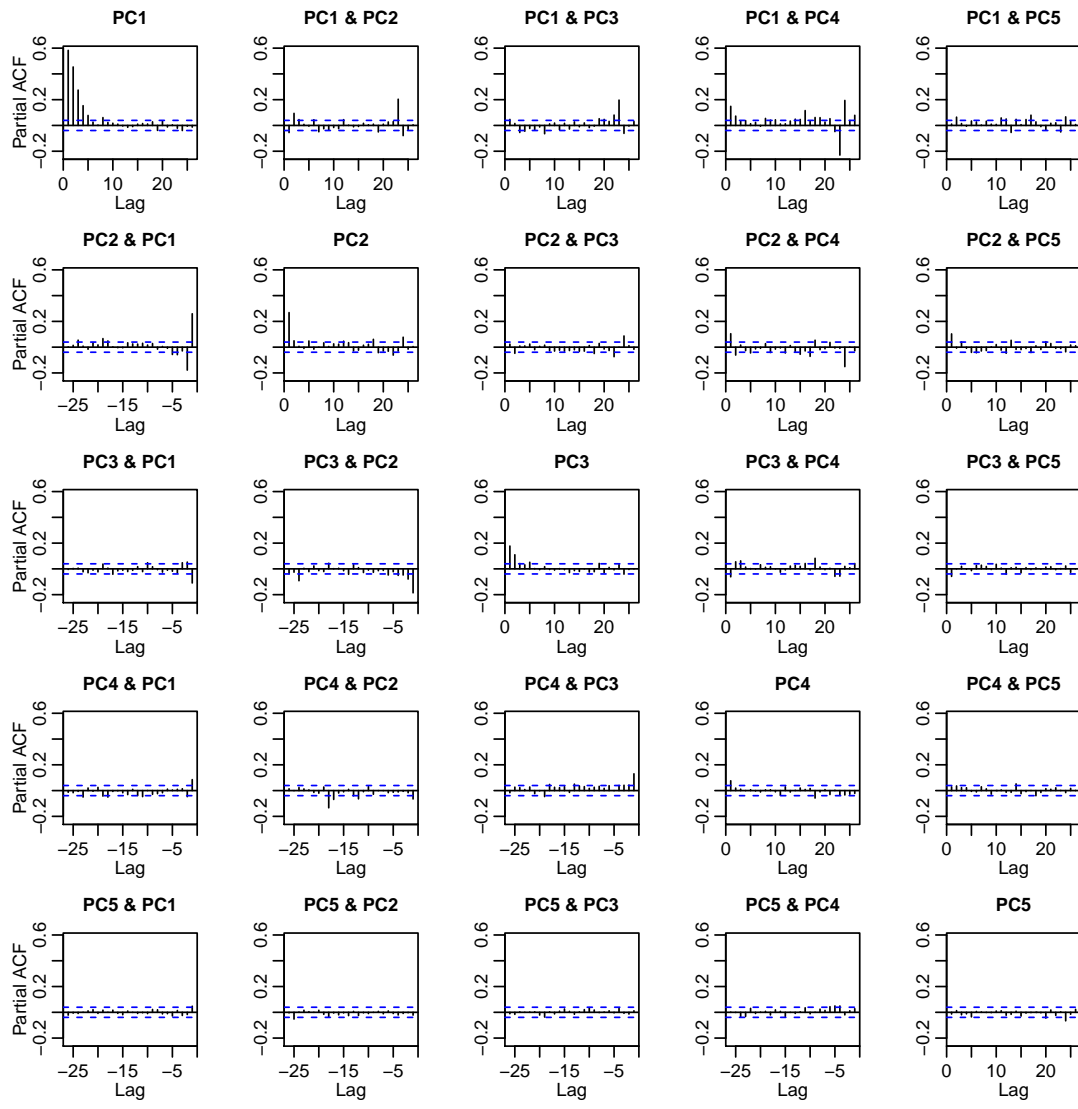


Figure 8-3: The partial autocorrelation of Bank of America’s principal components scores of daily-normalized liquidity measures.

while the R^2 for components four and five are significantly smaller at 0.08 and 0.01 respectively.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.3052	0.0215	14.21	0.0000 ***
PC2.11	-0.4264	0.0327	-13.03	0.0000 ***
PC3.11	0.2401	0.0302	7.96	0.0000 ***
PC4.11	-0.0417	0.0341	-1.22	0.2207
PC5.11	-0.0859	0.0337	-2.55	0.0109 *
PC1.12	0.4739	0.0254	18.65	0.0000 ***
PC2.12	0.1007	0.0270	3.73	0.0002 ***
PC3.12	0.0121	0.0278	0.44	0.6635
PC4.12	0.0712	0.0327	2.18	0.0297 *
PC5.12	0.0699	0.0335	2.09	0.0371 *
const	0.0061	0.0329	0.19	0.8518

Table 8.1: The first principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.3764	0.0145	26.00	0.0000 ***
PC2.11	0.3961	0.0221	17.96	0.0000 ***
PC3.11	-0.0330	0.0203	-1.62	0.1045
PC4.11	0.1935	0.0230	8.43	0.0000 ***
PC5.11	0.1467	0.0227	6.46	0.0000 ***
PC1.12	-0.1903	0.0171	-11.12	0.0000 ***
PC2.12	0.0445	0.0182	2.45	0.0144 *
PC3.12	-0.0459	0.0187	-2.46	0.0141 *
PC4.12	-0.0631	0.0221	-2.86	0.0043 **
PC5.12	-0.0089	0.0226	-0.39	0.6931
const	0.0022	0.0221	0.10	0.9201

Table 8.2: The second principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	-0.1493	0.0149	-10.02	0.0000 ***
PC2.11	-0.2123	0.0227	-9.35	0.0000 ***
PC3.11	0.1662	0.0209	7.94	0.0000 ***
PC4.11	-0.1148	0.0236	-4.86	0.0000 ***
PC5.11	-0.0763	0.0234	-3.26	0.0011 **
PC1.12	0.0573	0.0176	3.25	0.0012 **
PC2.12	-0.0768	0.0187	-4.10	0.0000 ***
PC3.12	0.1094	0.0193	5.68	0.0000 ***
PC4.12	0.0574	0.0227	2.53	0.0115 *
PC5.12	0.0011	0.0232	0.05	0.9630
const	-0.0018	0.0228	-0.08	0.9377

Table 8.3: The third principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.1117	0.0130	8.56	0.0000 ***
PC2.11	-0.0318	0.0199	-1.60	0.1100
PC3.11	0.1042	0.0183	5.69	0.0000 ***
PC4.11	0.0813	0.0207	3.93	0.0001 ***
PC5.11	0.0384	0.0205	1.88	0.0606
PC1.12	-0.0485	0.0154	-3.14	0.0017 **
PC2.12	-0.0165	0.0164	-1.00	0.3157
PC3.12	0.0438	0.0169	2.60	0.0094 **
PC4.12	0.0207	0.0199	1.04	0.2974
PC5.12	0.0360	0.0203	1.77	0.0772
const	0.0018	0.0199	0.09	0.9298

Table 8.4: The fourth principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.0577	0.0129	4.47	0.0000 ***
PC2.11	-0.0098	0.0197	-0.50	0.6176
PC3.11	0.0079	0.0181	0.43	0.6647
PC4.11	0.0347	0.0205	1.70	0.0901
PC5.11	-0.0093	0.0203	-0.46	0.6461
PC1.12	-0.0179	0.0153	-1.17	0.2411
PC2.12	-0.0034	0.0162	-0.21	0.8324
PC3.12	-0.0132	0.0167	-0.79	0.4289
PC4.12	0.0123	0.0197	0.62	0.5332
PC5.12	0.0130	0.0201	0.65	0.5182
const	-0.0011	0.0197	-0.06	0.9549

Table 8.5: The fifth principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

8.2 Citigroup Time Series Analysis

As with Bank of America, the scores plots, shown in Figure 8-4 indicate a very clear trend of seasonality across all of the components. While their zero means are contrived, depending on the component, the shape for each day seems to be a U-shape or in the case of component two, an inverted U-shape. The range of variation is greatest for component one, but it also appears to be the most consistent shape and height out of the components. Component two's peaks are not quite as smooth. Components two and three exhibit close to the same range while the four component shows significantly less time series structure but does vary within a smaller range of values.




Figure 8-4: The cumulative sum of Citigroup's principal components scores over time.

Figure 8-5: The first principal component's results from a second-order VAR model applied to Bank of America's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

8.2.1 Time Series Structure

The intuitions about which components have strong time series structure from Figure 8-4 can be solidified through the autocorrelations and cross-correlations of the components' scores. While all components have statistically significant coefficients for autocorrelation, these are strongest for component one, rather weak for components two and three, and then close to non-significant for component four. From the cross-correlations, we can see that component two and component three values influence component one values for the next five to fifty minutes. Component five slightly influences component one throughout all of the lags, but not to a very large extent.

The autocorrelations in Figure 8-7 indicate that most components are first order with a few of the components themselves being second or third order. Beyond these, for component one, its partial correlations is significant such that a VAR(6) could potentially be fit to just that component. However, we apply the Schwarz Criterion to the scores, which results in a VAR(3) model.

8.2.2 Vector Autoregressive Model Regressive Model

The results of the VAR model estimation are below. The component R^2 values are 0.3685, 0.416, 0.1849, and 0.1047 for components one, two, three, and four respectively. The first two measures are fit very well, while it tapers off with later components. While the time dependencies of the first component, which represents the volume, depth, and trade count measures, is somewhat expected, as these measures all demonstrated significant intraday liquidity patterns, the extent to which the VAR model describes component two, is somewhat surprising. Component two describes a weighted average of the proportional quoted spread, the proportional effective spread, and composite liquidity. While none of these measures seemed to exhibit clear intraday patterns individually, together in component two they have high time interdependencies and predictability as demonstrated by its high R value of roughly 0.64. This suggests that component two may capture a significant temporal

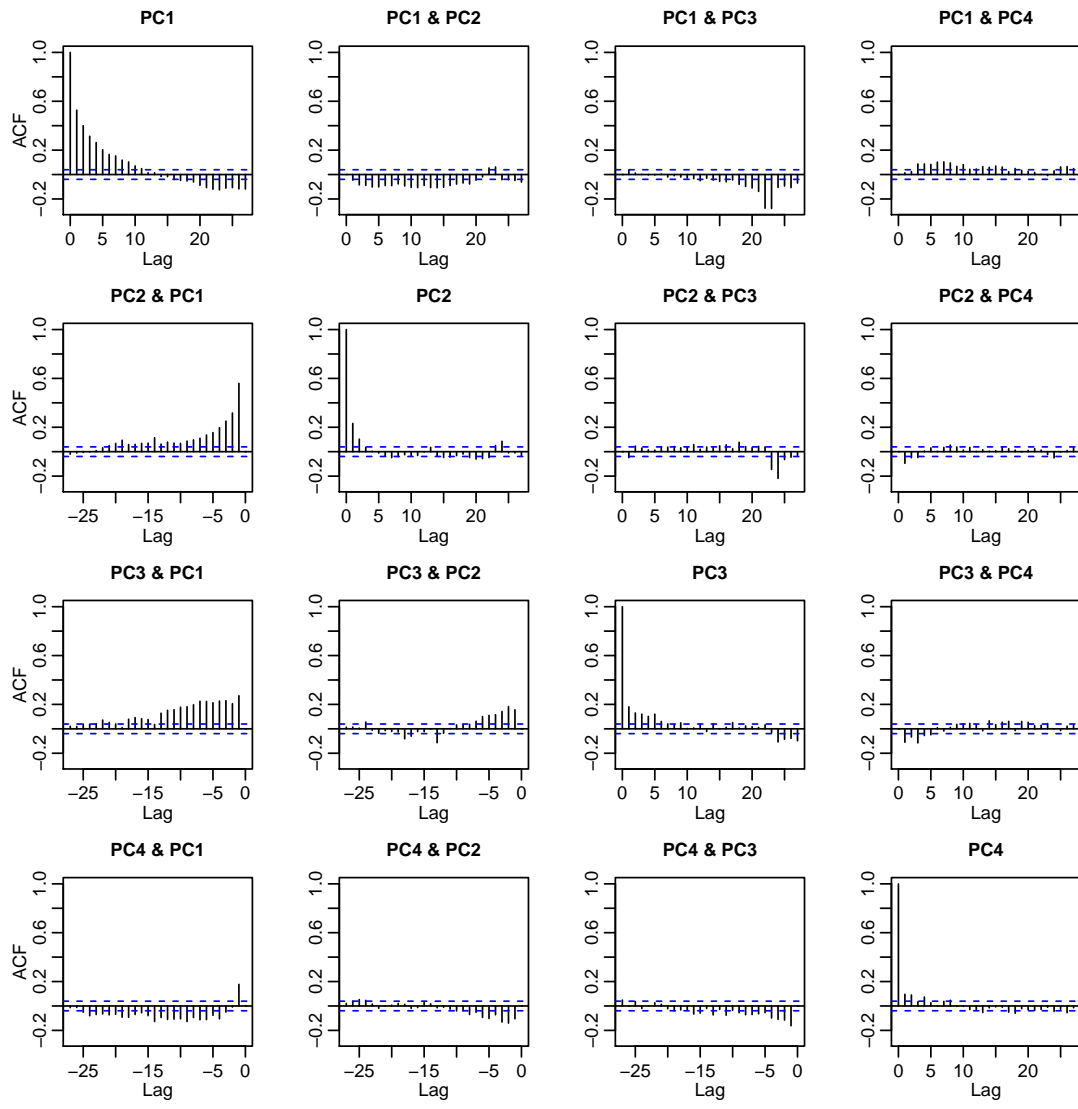


Figure 8-6: The autocorrelation and cross-correlation of Citigroup's principal components scores of daily-normalized liquidity measures.

dimension of liquidity that each measure alone is unable to.

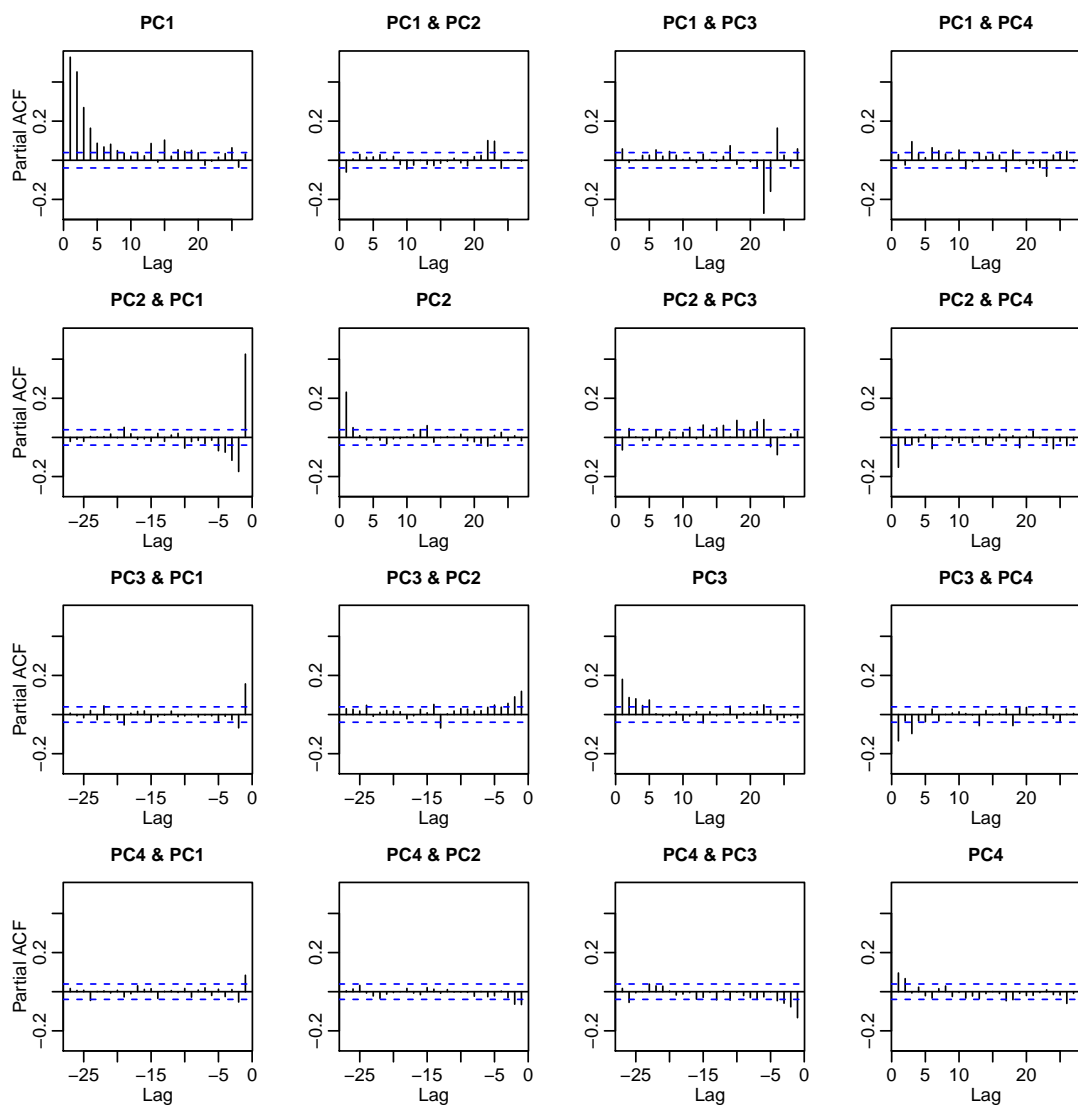


Figure 8-7: The partial autocorrelation of Citigroup's principal components scores of daily-normalized liquidity measures.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.2335	0.0232	10.05	0.0000 ***
PC2.11	-0.4004	0.0336	-11.90	0.0000 ***
PC3.11	-0.1529	0.0339	-4.51	0.0000 ***
PC4.11	-0.1490	0.0363	-4.10	0.0000 ***
PC1.12	0.3352	0.0346	9.69	0.0000 ***
PC2.12	-0.2119	0.0356	-5.95	0.0000 ***
PC3.12	-0.1441	0.0338	-4.26	0.0000 ***
PC4.12	-0.1313	0.0366	-3.58	0.0003 ***
PC1.13	0.2814	0.0322	8.73	0.0000 ***
PC2.13	0.0364	0.0234	1.56	0.1199
PC3.13	0.0041	0.0289	0.14	0.8872
PC4.13	0.0909	0.0347	2.62	0.0089 ***
const	0.0047	0.0345	0.14	0.8910

Table 8.6: The first principal component's results from a third-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.5410	0.0168	32.12	0.0000 ***
PC2.11	0.3424	0.0244	14.05	0.0000 ***
PC3.11	-0.0121	0.0246	-0.49	0.6232
PC4.11	-0.0613	0.0263	-2.33	0.0200 *
PC1.12	-0.1177	0.0251	-4.69	0.0000 ***
PC2.12	0.1439	0.0258	5.58	0.0000 ***
PC3.12	0.0999	0.0245	4.07	0.0000 ***
PC4.12	0.0137	0.0265	0.52	0.6058
PC1.13	-0.1233	0.0234	-5.28	0.0000 ***
PC2.13	0.0059	0.0169	0.35	0.7268
PC3.13	0.0016	0.0210	0.08	0.9390
PC4.13	-0.0351	0.0252	-1.40	0.1630
const	0.0032	0.0250	0.13	0.8992

Table 8.7: The second principal component's results from a third-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.1943	0.0152	12.75	0.0000 ***
PC2.11	0.1239	0.0220	5.62	0.0000 ***
PC3.11	0.1437	0.0222	6.46	0.0000 ***
PC4.11	-0.0567	0.0238	-2.38	0.0174 *
PC1.12	-0.0386	0.0227	-1.70	0.0888
PC2.12	0.1050	0.0233	4.50	0.0000 ***
PC3.12	0.0707	0.0222	3.19	0.0015 **
PC4.12	0.0022	0.0240	0.09	0.9283
PC1.13	-0.0270	0.0211	-1.28	0.2007
PC2.13	0.0572	0.0153	3.73	0.0002 ***
PC3.13	0.0806	0.0190	4.25	0.0000 ***
PC4.13	-0.0972	0.0228	-4.27	0.0000 ***
const	0.0016	0.0226	0.07	0.9427

Table 8.8: The third principal component's results from a third-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.1082	0.0131	8.25	0.0000 ***
PC2.11	-0.0063	0.0190	-0.33	0.7383
PC3.11	-0.0634	0.0191	-3.31	0.0009 ***
PC4.11	0.0725	0.0205	3.53	0.0004 ***
PC1.12	-0.0599	0.0195	-3.07	0.0022 **
PC2.12	-0.0704	0.0201	-3.51	0.0005 ***
PC3.12	-0.0691	0.0191	-3.62	0.0003 ***
PC4.12	0.0503	0.0207	2.43	0.0151 *
PC1.13	0.0106	0.0182	0.58	0.5615
PC2.13	-0.0303	0.0132	-2.30	0.0215 *
PC3.13	-0.0583	0.0163	-3.57	0.0004 ***
PC4.13	-0.0066	0.0196	-0.34	0.7361
const	0.0012	0.0194	0.06	0.9506

Table 8.9: The fourth principal component's results from a third-order VAR model applied to Citigroup's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

8.3 Financial Sector ETF Time Series Analysis

As seen with Bank of America and Citigroup, the scores plots of the Financial Sector ETF, found in Figure 8-8, show strong intraday patterns as indicated by the specific shape repeated five times every week. However, this pattern starts to become less clear by the third and fourth measures, and it not very apparent at all in the fifth component score plot.

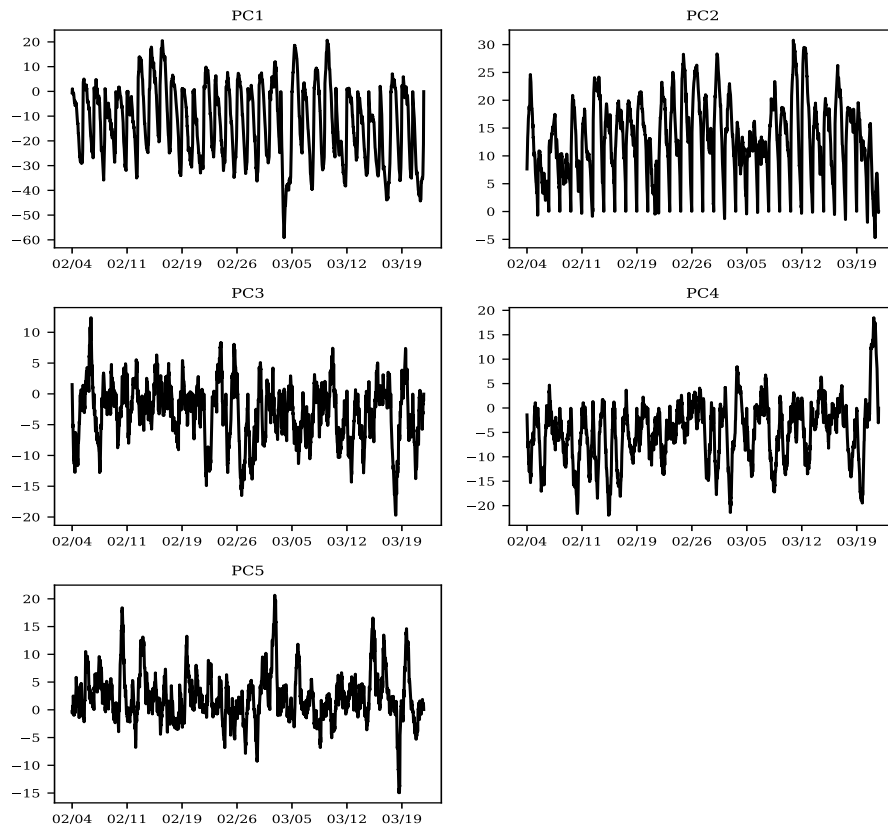


Figure 8-8: The cumulative sum of the Financial Sector ETF's principal components scores over time.

8.3.1 Time Series Structure

These intuitions are further confirmed in the autocorrelations and cross-correlations in Figure 8-9. Components one, two, and four to a lesser extent have meaningful time series structure, while there is effectively no significant autocorrelation for

components three and five. There is cross-correlation where component two values precede component one values five to twenty-five minutes before. There is some cross correlation between components one and three and one and four, but these correlation coefficients are very low.

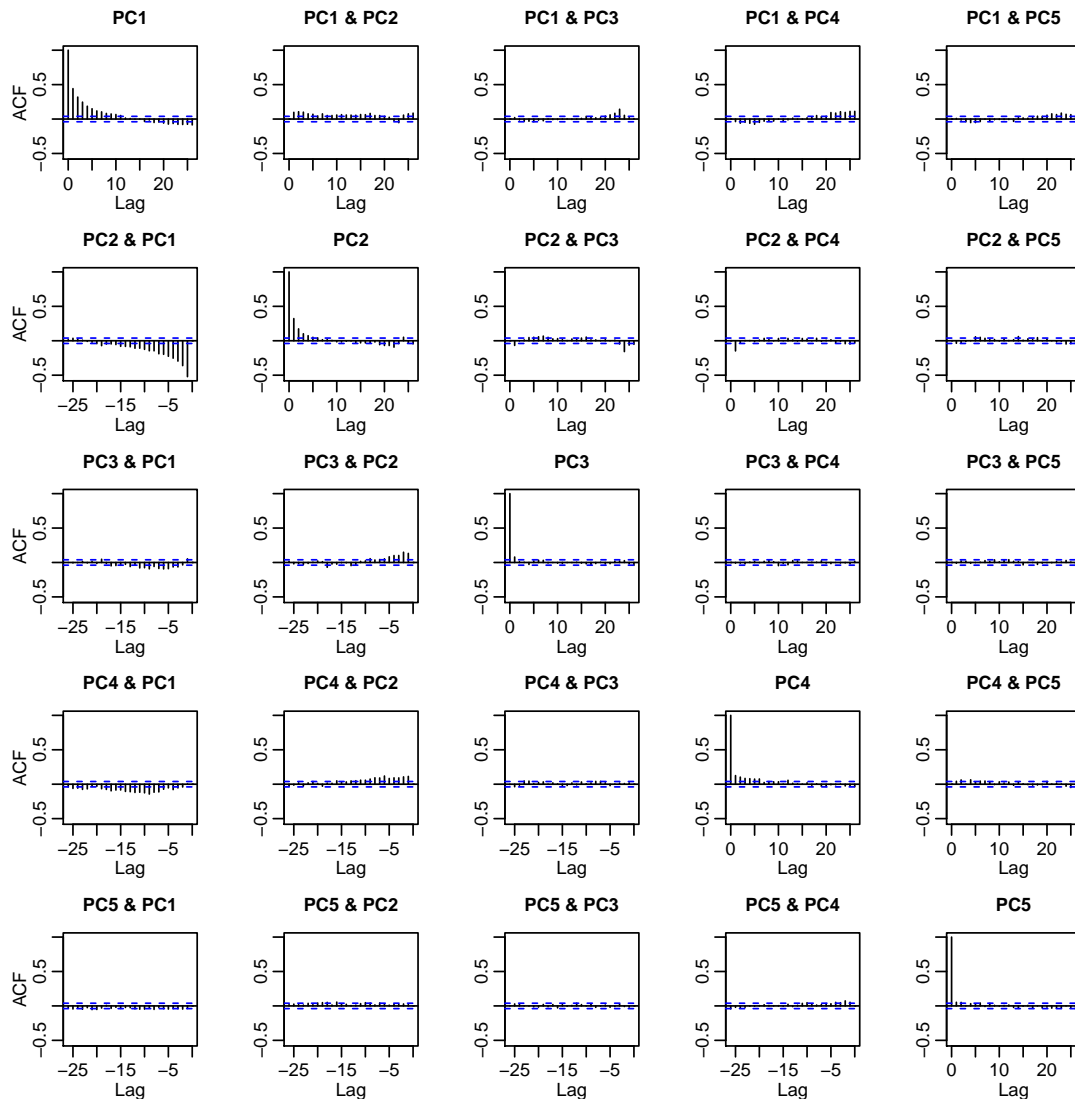


Figure 8-9: The autocorrelation and cross-correlation of the Financial Sector ETF's principal components scores of daily-normalized liquidity measures.

The partial autocorrelations shown in Figure 8-10 indicates three to four-order VAR models for liquidity components one, two, and four, but aside from these the

components partial autocorrelation suggest a lower order model may be more appropriate.

8.3.2 Vector Autoregressive Model

The Schwarz Criterion dictates using a second-order VAR(2) model to fit the data. The results are shown below. To summarize the results in term of R^2 , in order of components, their values are 0.2706, 0.4425, 0.03724, 0.04822, and 0.01515. As with Citigroup, the first two components appear to be modeled quite well by the VAR model as indicated by their large R^2 values. Component one is the volume, depth, and trade count component we have observed countless times in previous analyses; it is a component which has a clearly demonstrated temporal dependency across all securities and time periods of analysis.

However, component two's large R^2 is perhaps unexpected. Component two describes a weighted average of composite liquidity and the Martin Index contrasted against the quoted spread. It is a measure which perhaps encompasses the price change between consecutive trades and how the dynamics between quoted spread and depth respond to this compared to the way in which the quoted spread alone responds. Because none of these measures individually possessed strong intraday time dependencies, the large R value of component two suggests that their combination describes a meaningful aspect of intraday liquidity, one that is not captured by the components alone and one that contains significant time structure. Given this, it is perhaps a measure that could be particularly useful in forecasting models for its relative predictability.

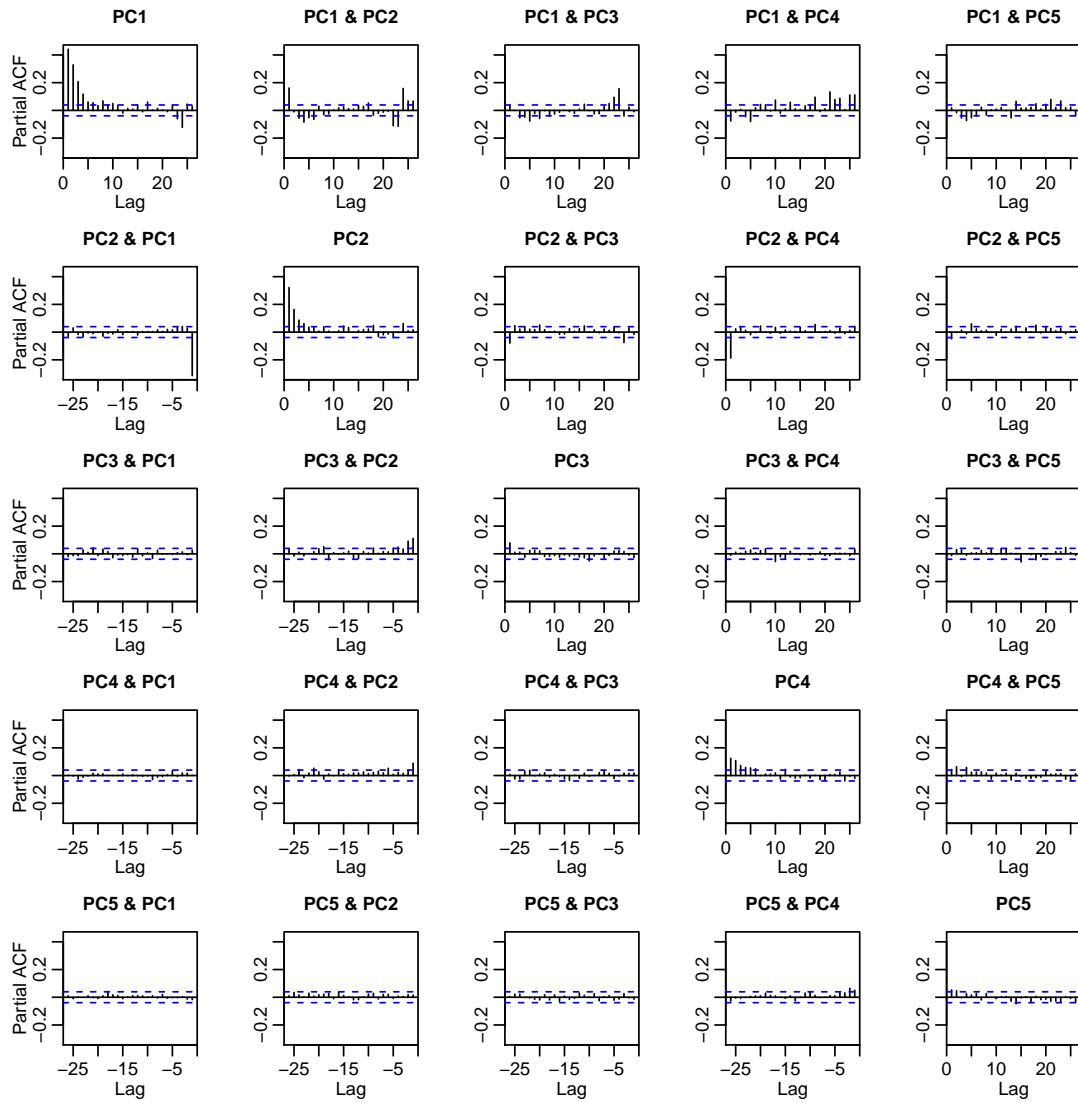


Figure 8-10: The partial autocorrelation of the Financial Sector ETF's principal components scores of daily-normalized liquidity measures.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.2973	0.0203	14.65	0.0000 ***
PC2.11	0.4754	0.0392	12.13	0.0000 ***
PC3.11	0.0087	0.0342	0.25	0.8000
PC4.11	-0.0672	0.0368	-1.83	0.0680
PC5.11	0.0489	0.0374	1.31	0.1908
PC1.12	0.3473	0.0246	14.09	0.0000 ***
PC2.12	-0.0202	0.0314	-0.64	0.5196
PC3.12	0.0014	0.0340	0.04	0.9667
PC4.12	-0.0029	0.0372	-0.08	0.9373
PC5.12	-0.0126	0.0374	-0.34	0.7360
const	-0.0029	0.0372	-0.08	0.9370

Table 8.10: The first principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	-0.3418	0.0106	-32.21	0.0000 ***
PC2.11	0.2994	0.0205	14.61	0.0000 ***
PC3.11	-0.1139	0.0179	-6.37	0.0000 ***
PC4.11	-0.2198	0.0192	-11.42	0.0000 ***
PC5.11	-0.0513	0.0196	-2.63	0.0087 **
PC1.12	0.0274	0.0129	2.12	0.0339 *
PC2.12	0.1642	0.0164	10.01	0.0000 ***
PC3.12	0.0489	0.0178	2.75	0.0060 **
PC4.12	0.0234	0.0195	1.20	0.2291
PC5.12	-0.0035	0.0195	-0.18	0.8563
const	-0.0009	0.0195	-0.05	0.9617

Table 8.11: The second principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	0.0245	0.0119	2.07	0.0389 *
PC2.11	0.0768	0.0229	3.35	0.0008 ***
PC3.11	0.0655	0.0200	3.28	0.0011 **
PC4.11	-0.0265	0.0215	-1.23	0.2184
PC5.11	-0.0169	0.0218	-0.77	0.4387
PC1.12	-0.0067	0.0144	-0.46	0.6441
PC2.12	0.0923	0.0183	5.03	0.0000 ***
PC3.12	0.0106	0.0199	0.54	0.5923
PC4.12	0.0114	0.0218	0.52	0.6003
PC5.12	0.0304	0.0218	1.39	0.1644
const	0.0001	0.0218	0.01	0.9958

Table 8.12: The third principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	-0.0065	0.0110	-0.59	0.5527
PC2.11	0.1064	0.0212	5.02	0.0000 ***
PC3.11	0.0100	0.0185	0.54	0.5887
PC4.11	0.1028	0.0199	5.16	0.0000 ***
PC5.11	0.0349	0.0202	1.73	0.0845
PC1.12	0.0179	0.0133	1.34	0.1788
PC2.12	0.0395	0.0170	2.33	0.0201 *
PC3.12	0.0155	0.0184	0.84	0.4006
PC4.12	0.1083	0.0202	5.38	0.0000 ***
PC5.12	0.0633	0.0202	3.13	0.0018 **
const	-0.0011	0.0202	-0.05	0.9580

Table 8.13: The fourth principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

	Estimate	Std. Error	t value	Pr(> t)
PC1.11	-0.0116	0.0109	-1.07	0.2855
PC2.11	0.0078	0.0210	0.37	0.7115
PC3.11	-0.0153	0.0183	-0.84	0.4017
PC4.11	0.0387	0.0197	1.97	0.0494 *
PC5.11	0.0450	0.0200	2.25	0.0246 *
PC1.12	-0.0125	0.0132	-0.95	0.3433
PC2.12	0.0188	0.0168	1.12	0.2621
PC3.12	-0.0052	0.0182	-0.28	0.7770
PC4.12	0.0644	0.0199	3.23	0.0012 **
PC5.12	0.0473	0.0200	2.37	0.0180 *
const	0.0003	0.0199	0.02	0.9867

Table 8.14: The fifth principal component's results from a second-order VAR model applied to the Financial Sector ETF's principal component scores. * indicates t value is less than 0.001. ** indicates t value is less than 0.01. *** indicates t value is less than 0.05.

Chapter 9

Conclusion

In summary, intraday data was collected for Bank of America, Citigroup, and the Financial Sector ETF over a 32 trading-day period from February 4, 2019 to March 20, 2019. From this trade and quote data, we construct 2,496 five-minute aggregate time bars for each security and calculate a series of spread, volume, depth, trade count, and price change liquidity measures. We examine the summary statistics of these liquidity measures to find that, as expected, they present as complex mixtures with significant skewness and/or kurtosis. Furthermore, we aggregate daily-normalized liquidity measures into 15-minute time bars and demonstrate that there are strong patterns of systematic intraday liquidity, the most noticeable of which are in volume, depth, and trade count measures.

We perform principal components analysis on the liquidity measures through which we identify key liquidity dimensions in each security and common liquidity dimensions across them. We then aim to understand the time structure of these principal component scores through the use of vector autoregressive models. Finally, the same methodologies of principal components analysis and time series analysis are applied to daily-normalized liquidity measures in order to better understand the intraday, rather than multi-day, dynamics of liquidity. In classical time series analysis, it is very common to transform the time series to eliminate non-stationarity by taking differences of the series. These analyses do not follow this approach because

the analysis of differences can obscure long-term time dependence and variability. The analysis is concerned with understanding variation and interdependence in the measures on both an intraday basis and a multi-day basis spanning several weeks. These methods suggest useful initial models of the dynamics of these measures. Future work would analyze more stocks and understand the degree to which there are different liquidity profiles and models for different stocks and sectors. The results of the research presented here can help guide the development of such models in futures research.

Bibliography

- [1] Kenneth Ahern. Do proxies for informed trading measure informed trading? evidence from illegal insider trades. 2018.
- [2] Yakov Amihud. Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1):3156, 2002.
- [3] Yakov Amihud, Haim Mendelson, and Lasse Heje Pedersen. Asset pricing and the bidask spread. *Market Liquidity*, page 946, Apr 1986.
- [4] Michael J. Barclay and Jerold B. Warner. Stealth trading and volatility. *Journal of Financial Economics*, 34(3):281305, 1993.
- [5] A. Brunner. Measurement concepts for liquidity on securities markets. *Center for Financial Studies*, 1996.
- [6] Charles Cao, Oliver Hansch, and Xiaoxin Wang Beardsley. The informational content of an open limit order book. *SSRN Electronic Journal*, 2004.
- [7] Tarun Chordia, Richard W. Roll, and Avanidhar Subrahmanyam. Commonality in liquidity. *SSRN Electronic Journal*, 1999.
- [8] Tarun Chordia, Richard W. Roll, and Avanidhar Subrahmanyam. Market liquidity and trading activity. *SSRN Electronic Journal*, 2000.
- [9] David Easley, Marcos M. Lopez De Prado, and Maureen Ohara. Vpin and the flash crash: A comment. *SSRN Electronic Journal*, 2012.
- [10] Rachel Evans and Carolina Wilson. How etfs became the market.
- [11] Ruslan Goyenko, Craig W. Holden, and Charles Trzcinka. Do measures of liquidity measure liquidity? *SSRN Electronic Journal*, 2008.
- [12] Yasushi Hamao and Joel Hasbrouck. Securities trading in the absence of dealers: Trades and quotes on the tokyo stock exchange. *Review of Financial Studies*, 8(3):849878, 1995.
- [13] Lawrence E. Harris. Minimum price variations, discrete bidask spreads, and quotation sizes. *Review of Financial Studies*, 7(1):149178, 1994.

- [14] Joel Hasbrouck. Measuring the information content of stock trades. *The Journal of Finance*, 46(1):179, 1991.
- [15] Joel Hasbrouck. Trading costs and returns for us equities: The evidence from daily data. *SSRN Electronic Journal*, 2003.
- [16] Doron Israeli, Charles M.c. Lee, and Suhas A. Sridharan. Is there a dark side to exchange traded funds (etfs)? an information perspective. *SSRN Electronic Journal*, 2015.
- [17] Charles M. Jones, Gautam Kaul, and Marc L. Lipson. Information, trading, and volatility. *Journal of Financial Economics*, 36(1):127154, 1994.
- [18] Frank De Jong and Barbara Rindi. The microstructure of financial markets. 2009.
- [19] Thomas H. Mcinish and Robert A. Wood. An analysis of intraday patterns in bid/ask spreads for nyse stocks. *The Journal of Finance*, 47(2):753, 1992.
- [20] Maureen Ohara. Presidential address: Liquidity and price discovery. *The Journal of Finance*, 58(4):13351354, 2003.
- [21] Richard Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance*, 39(4):1127, 1984.
- [22] Rico von. Wyss. *Measuring and predicting liquidity in the stock market*. PhD thesis, 2004.