A Time and Frequency Domain Analysis of Contrarian Trading Strategies

by

Shomesh E. Chaudhuri

B.S., Harvard University (2011)

Submitted to the Department of Electrical Engineering and Computer Science
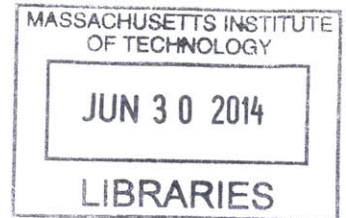in partial fulfillment of the requirements for the degree of

Master of Science in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2014

Signature redacted

Author . . . .
Department of Electrical Engineering and Computer Science
May 21, 2014

Signature redacted

Certified by. . . . .
Andrew W. Lo
Charles E. and Susan T. Harris Professor
Thesis Supervisor

Signature redacted

Accepted by . . . .
Leslie A. Kolodziejski
Chair, Committee on Graduate Students

A Time and Frequency Domain Analysis of Contrarian Trading Strategies

by

Shomesh E. Chaudhuri

Abstract

This thesis applies time and frequency domain analyses to a high-frequency market making strategy to study the profitability of liquidity provision over multiple time horizons from 1964 to 2013. Using daily returns and inside quotes, we provide evidence that widening spreads on the NASDAQ National Market System in the late 1980s and early 1990s were facilitated by implicit pricing agreements amongst security dealers. In contrast, we show that regulatory changes, such as decimalization, and the development of liquidity providing algorithmic trading strategies acted to narrow spreads and reduce transaction costs. Increasing the focus of our analytical lens to the intraday level we find that, over the past two decades, market making profitability has been higher and sensitivity to market turbulence has been lower at shorter time horizons providing strong incentives for traders to move into higher frequency domains. As an informative example we show that, during the 2010 Flash Crash, high-frequency realizations of our market making strategy are unaffected by the rapid market downturn while slower realizations are caught on the wrong side of markets as prices fall. This last observation leads us to consider a frequency decomposition of average portfolio returns to characterize the profitability of our trading strategy at different time horizons. We use this spectral technique to demonstrate that turbulence or momentum in one frequency band may substantially affect the profitability of contrarian trading strategies that target a given time horizon while leaving other strategies relatively unaffected.

Acknowledgments

Thank you to my family, teachers and friends without whom none of this would be possible. I cannot express in words how grateful I am to those who have supported and continue to support me in my personal, academic and professional life.

In particular, I would like to thank Professor Andrew Lo for providing guidance, wisdom and encouragement throughout this entire process. I would also like to thank the LFE administrators, Patsy and Jayna, for helping me connect with the resources needed for my research. Finally, I would like to thank my best friend Meaghan, my parents Janet and Sujeet, and my brother Sheel for all the love, support, motivation and wisdom they have given me.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The past couple of decades have been witness to some of the highest periods of volatility in the history of the stock market. This increase in volatility has incidentally occurred during a major technological and environmental shift in the financial industry leading many to wonder if and how these effects are related. The advent of electronic communication networks (ECNs) – electronic systems that facilitate trading – has been linked to a steep rise in overall trading volume, causing some to speculate that algorithmic trading firms could be the driving force behind the unusually high market volatility. Shorter investment horizons, higher turnover and an increase in the complexity of interacting algorithms may well pose a risk to market stability. On the other hand, proponents of algorithmic trading argue that lower latency times and more intelligent trading strategies have increased market liquidity and efficiency, ultimately acting to mitigate market volatility. Unfortunately, the problem of determining precisely how technology and new regulations have affected market conditions has proven difficult, especially in the wake of the dot-com and housing market bubbles.

This paper will apply time and frequency domain techniques to help untangle the role technology and regulations have played in shaping recent market events. The scope of the analysis, ranging from 1964 to 2013, complemented with the laser-like precision of the mathematical tools employed will enable us to uncover subtle yet important market shifts. While the analytical tools used in this paper are based

on very simple trading strategies that are not representative of the statistical arbitrage industry as a whole, and a simulation of a more sophisticated set of strategies would undoubtedly paint a more complex picture, our goal is to expose some of the fundamental relationships that exist in the marketplace today. Understanding the connection between technology and market stability has become especially important as recent events like the 2010 Flash Crash have led many to question the potential negative impacts algorithmic trading could have on markets. Only by understanding the interactions of all financial activity will we be able to implement well-designed regulations that curb the risk of market instability.

1.1   Literature Review

We begin by studying the impact the bid-ask bounce, market liquidity and investor overreaction have on price reversals at the one-day time horizon. De Bondt and Thaler (1985) find that when stocks are grouped into winners and losers by their past three- to five- year returns, the past losers tend to outperform past winners in the future. They attribute these long-term reversals to irrational investor overreaction to past information caused by behavioral decision theory biases. Since then, many other return anomalies across a spectrum of time horizons have been studied and their causes and significance debated.

At the daily returns resolution, one method the overreaction hypothesis has been studied involves analyzing stock returns immediately following large price changes. Brown, Harlow and Ticnic (1988) find significant reversals after large price declines and provide a rational behavior explanation, the uncertain information hypothesis. They suggest that in addition to increasing measurable risk, a noisy piece of favorable or unfavorable news causes risk-averse investors to set stock prices below their conditional expected value. As the uncertainty is resolved, subsequent price changes will tend to be positive on average. Restricting event studies to one-day price declines of at least 10 percent, both Bremer and Sweeney (1991) and Cox and Peterson (1994) also find significant price reversals. Cox and Peterson suggest that much of this price

reversal can be explained by the bid-ask bounce, and similarly Atkins and Dyl (1990) find that price reversals following large single day returns are small relative to the bid-ask spread.

The mechanism for these bid-ask bounce price reversals works as follows. One-day price declines (increases) are associated with selling (buying) pressure, increasing the probability that the closing price associated with the last transaction is at a the bid (ask) price. The larger the magnitude of the price decrease (increase), the larger the probability that the price is associated with the bid (ask). If, on the next day, selling and buying pressures re-equilibrate, then the price will have an equal chance of closing at either the bid or the ask and will, on average, induce a negative autocorrelation in returns. Moreover, since the inventory risks borne by market makers increase during large changes in price, spreads may widen in response to thinning liquidity, amplifying the size of the subsequent bid-ask bounce.

For our daily returns analysis, the availability of closing transaction prices as well as inside bid and ask quotes for the NASDAQ National Market System (NMS) securities allows us to study short-term reversals with respect to the bid-ask spread. If the bid-ask bounce plays an important role in price reversal, then we should expect that the profits of a simulated contrarian trading strategy designed to take advantage of negative autocorrelation should 1) be correlated with the size of the percentage bid-ask spread, and 2) diminish if prices are calculated as the midpoint between the inside bid and ask quotes. We would also expect these profits to decrease over time as markets become more liquid due to increased market making competition driven by advances in alternative trading systems like ECNs.

The adoption of computer technology as well as the implementation of regulatory changes designed to promote greater transparency and competition have all helped electronic and algorithmic trading gain popularity over the past couple of decades. One example of this trend occurred in 1993, when the Professional Trader Rule, a regulation that barred the high-frequency traders of the day – SOES bandits – from trading on the NASDAQ's Small Order Execution System (SOES), was removed. These professional day traders, who only held positions for a few minutes, were able

15

to capitalize on short-term momentum and reversals by having their orders executed rapidly through ECNs like SelectNet and Instinet. By establishing positions before the slower market makers could update their quotes, SOES bandits were able trade within the bid-ask spread. This strategy earned a small average profit per trade, and would be executed thousands of times per month (Harris and Schultz 1998).

Harris and Schultz (1998) study the source of profits of SOES bandits and find that the ECNs SelectNet and Instinct were essential to making their strategies profitable. Their complementary paper (1997) examines the effects of a rule change in January 1994 that reduced the maximum SOES trade from 1000 to 500 shares and finds little evidence that quoted spreads were subsequently changed. Battalio, Hatch and Jennings (1997) study the impact SOES bandit trading had on market volatility. As with today's high-frequency traders, there was concern that SOES bandits, if pursuing a momentum strategy, could push a stock's price higher when it had been increasing and depress the price further if it was decreasing. They report that bandits did not in fact increase volatility, but rather facilitated faster price changes.

Over time, algorithmic trading grew in popularity as the trading environment continued to evolve in its favor. The reduction of the minimum tick size first from 1/8th to 1/16th of a dollar in 1997 and then to one cent in 2001 spurred innovation in the liquidity provision algorithmic trading sector, as smaller spreads required more sophisticated techniques to capture profits. The introduction of the Regulation National Market System (Reg. NMS) in 2007, designed to promote transparency by requiring trade orders to be posted nationally, further increased the premium for speed as traders could now profit from momentary price lags between exchanges (Capgemeni 2012). Catalyzed by advancements in fiber optics and digital processors, high-frequency traders (HFTs) have profoundly changed market microstructures. Today, despite representing only 2% of all trading firms, high-frequency trading (HFT) in the US accounts for almost 75% of all trades, and the average time a stock is held has been reduced to 22 seconds (Solman 2012). Looking ahead, the prevalence of HFT is expected to continue to grow as hedge funds and proprietary trading desks look to implement these highly profitable trading strategies globally and across various asset

classes.

A recent paper by Baron, Brogaard and Kirilenko (2012) examined the profitability of HFTs and found that an important determinant of a firm's profitability was its level of liquid provision. Aggressive liquid-taking HFTs, defined as HFT firms that initiate at least 60% of their trades, were found to be more profitable than liquid-providing passive HFTs. Similar to some of the literature on hedge funds (Jagannathan, Malakhov and Novikov 2010) but unlike mutual funds (Carhart 1997), they find that, as a group, HFTs consistently outperform the market.

Of the commonly deployed electronic trading strategies, this paper will focus on a form of liquidity providing statistical arbitrage. We simulate the performance of a short-term mean reversion strategy developed by Lo and MacKinlay (1990) that involves holding hundreds of securities for very short holding periods (measured in seconds to days). In doing so we will be able to quantify the systematic decline in the profitability of long/short equity strategies due to technological advances, environmental changes such as decimalization, and increased market making competition. The simulated strategy functions by buying recent losers and shorting recent winners. Contrarian trading strategies such as these, in effect, stabilize price swings, lower volatility, and provide liquidity in the marketplace. Traditionally, market specialists and dealers, who were compensated through the bid-ask spread, have held this role of liquidity provider but, as spreads have narrowed, hedge funds and propriety trading desks employing contrarian-like strategies have stepped in to compete with market makers. If these strategies do in fact reduce volatility, we should be able to show that decreasing contrarian trading profitability, caused by increased competition fueled by technological advancement, corresponds to greater market efficiency, higher levels of liquidity and lower levels of volatility. Similarly shocks to the system that cause hedge funds to reduce their risk exposure and withdraw capital from market making strategies will be marked by periods of increased CTS profitability, wider bid-ask spreads and higher levels of overall volatility.

If stocks exhibit momentum from sustained buy or sell pressure, then our mean reversion trading strategy will prove to be unprofitable. Khandani and Lo (2007)

leverage this fact and simulate a contrarian trading strategy to provide evidence that unprecedented losses suffered by quantitative hedge funds in the second week of August were triggered by the rapid deleveraging of one or more equity market-neutral portfolios. In their 2011 companion paper, they use 5-minute price level data to precisely identify the time and location of the unwinds that triggered this "Quant Meltdown". In a similar spirit, this paper aims to simulate the contrarian trading strategy to help uncover the role algorithmic trading, technology and regulation changes have played over the last five decades. The breadth, from 1964 to 2013, and depth, reaching a one-second price level resolution, of this analysis is unprecedented.

Finally, we modify a spectral analysis technique introduced by Engle (1974, 1978) and further developed by Hasbrouck and Sofianos (1993) to decompose the simulated contrarian trading strategy profits by frequency allowing us to determine in what time horizons these market making strategies earn profits. Our analysis shows that different trading strategies will be affected by different frequency components of return autocovariance, and therefore a diversification across frequencies can increase the risk-adjusted performance of an overall portfolio. For instance, Kirilenko et al. (2011) show that during the 2010 Flash Crash, HFTs, despite their large share of trading volume, appear to have operated at such a high speed that they were not affected by the price momentum or volatility. Intermediaries, on the other hand, traded on a slower timescale and seem to have been caught on the wrong side of the market as prices fell.

1.2   A Brief History of the National Market System

In the 1980s and early 1990s the NASDAQ experienced the introduction of real-time transaction reporting, the expansion of institutional trading systems including Instinet, increased institutional trading, the establishment and exploitation of the retail-oriented trading system SOES (Small Order Execution System), and market maker collusion to maintain artificially wide spreads. This section provides a brief introduction and history of these developments.

In February 1982, the National Association of Securities Dealers (NASD) established the National Market System (NMS) to provide more comprehensive, real-time information – high, low and last sale prices, cumulative volume figures, and bid and ask quotes – throughout the day for the most active stocks in the NASDAQ market. The new reporting rules significantly expanded transaction and price information available to NASD members, issuers and public investors in NASDAQ NMS securities. Market makers now had to report the actual price and number of shares in each transaction within 90 seconds as opposed to the nonreal-time reporting procedures[1] for non-NMS stocks. The NMS, which initially started with only 40 securities, had increased to 1180 of 4723 (25%) NASDAQ securities by the end of 1984 and accounted for 50% of total NASDAQ trading volume. By 1988, the number of NMS securities had grown to exceed 3000 and accounted for over 70% of total NASDAQ trading volume (Kothare and Laux 1995).

In 1984 the NASDAQ also experienced a rapid expansion in Instinet, a system designed to facilitate the direct trading of institutional stock positions, and consequently, an increase in the presence of institutional traders. In parallel, SOES, a system designed to allow individual investors to execute small trades quickly and give them access to similar orders as institutional traders, was developed. At the end of 1984, only 50 securities were listed on the pilot version of SOES.

In response to criticisms of difficulties in executing trades during the stock market crash of 1987, the NASD required dealers to automatically execute SOES trades via the NASDAQ computer system. Market makers were also prevented from having their quotes updated automatically, which provided an opportunity for short-term professional day traders (SOES bandits) to exploit cross-sectional discrepancies that arose from stale quotes using rapid buy and sell orders. In 1991, the NASD perceived this bandit trading as a significant problem to the public's trading costs, arguing that these trades distorted prices, increased volatility and caused dealers to widen spreads to compensate for their increased risk, thereby reducing liquidity. They approached the US Securities & Exchange Commission (SEC) to expand the definition

---

[1]last sales prices and minute-to-minute volume updates were not available.

of restricted trading on SOES to include more short-term professional traders. These amendments to the Professional Trader Rule were eventually dismissed in court as being vague and arbitrary. Overall however, SOES bandits restricted their trading to the largest and most active NMS stocks and only accounted for 2% of total NASDAQ volume in 1995 (Kothare and Laux 1995).

Kothare and Laux (1995) find little evidence that the abuse of SOES by day traders influenced transaction costs, and instead suggest that increased institutional activity widened bid-ask spreads. They report that bid-ask spreads for NMS stocks rose sharply (on the order of 100% overall) from the mid-1980s to mid-1990s and that growth in institutional trading may have created difficult to manage order flow which ultimately increased market maker inventory risk. They also suggest that if institutions were privately informed investors, then their growing presence would have increased dealers' exposure to asymmetric information risk, and consequently, broadened spreads.

On May 26 and 27, 1994 several national newspapers including the *Wall Street Journal* reported the evidence provided by Christie and Schultz (1994a) that market makers of active NASDAQ stocks were implicitly colluding to maintain wide spreads by avoiding odd-eighth quotes. In a follow up paper Christie and Schultz (1994b) show that on Friday May 27, 1994, the average inside spread for Amgen, Cisco Systems, and Microsoft fell to about $0.15 per share after varying between $0.25 and $0.45 during the preceding 17 months. A similar reduction in spreads was reported for Apple the following trading day. The authors find no changes in the cost of market making and attribute the decrease in transaction costs to a fundamental shift in the use of odd-eighth quotes and the breakdown of implicit price agreements between dealers. Market maker competition for order flow was supposed to drive down transaction costs leading to better execution for investors, however this finding undermines that premise and leads us to ask: how significant was collusion to increasing spreads throughout the 1980s and early 1990s? This paper provides evidence that in fact, market maker collusion and the avoidance of odd-eighth quotes was an important factor in widening spreads.

1.3 Dataset

In this thesis we use three sources of data for our analysis. Our daily returns analysis focuses on all stocks in the University of Chicago's CRSP Database from January 1, 1964 to December 31, 2013. Specifically, we only use US common stocks (CRSP share code 10 and 11), which eliminates REIT's, ADR's, and other types of securities, and we drop stocks with share prices below $5 and above $2,000. CRSP occasionally reports returns and prices that are based on bid-ask quote midpoints and, unless otherwise specified, these values are filtered out. Transaction prices and inside bid-ask quotes at each day's close[2] for NASDAQ NMS and Small Cap Market securities are not reported until November 1, 1982 and June 15, 1992, respectively. Prior to those dates, NASDAQ security prices are given as the bid-ask quote midpoints and the quotes themselves are not recorded. Furthermore, the bid quotes are missing for all NASDAQ NMS securities in February 1986.

All of our intraday returns analysis focuses on the members of the S&P Composite 1500, which includes all stocks in the S&P 500 (large-cap), S&P 400 (mid-cap) and S&P 600 (small-cap) indexes. Components of the S&P 1500 were based on memberships of the first day of the month and were obtained from Standard & Poor's Compustat database via the Wharton Research Data Services (WRDS) platform. To improve our temporal resolution, we use the NYSE Trade and Quote (TAQ) database. This database contains intraday transactions and quotes data for all securities listed on the NYSE, the American Stock Exchange, and the NASDAQ NMS and Small Cap Market. The dataset consists of the Daily National Best Bids and Offers (NBBO) File, the Daily Quotes File, the Daily TAQ Master File, and the Daily Trades File. For the purposes of this thesis, we only use actual trades as reported in the Daily Trades File. This file includes information such as the security symbol, trade time, size, exchange on which the trade took place, as well as a few condition and correction flags. We only use trades that occur during normal trading hours (9:30am to 4:00pm) and discard all records that have Trade Correction Indicator field entries other than

---

[2]with the close being at 4:00 PM ET

"00", which signifies a regular trade that was not corrected, changed or canceled. We also remove all trades that were reported late or out of sequence, according to the Sale Condition field (entries B, C, G, L, M, N, P, Q, R, S, T, W, Z and/or 9). These filters have been used in other studies based on TAQ data; see, for example Khandani and Lo (2011), Christie, Harris and Schultz (1994) or Chordia, Roll and Subrahmanyam (2001). See the TAQ documentation for further details.

Chapter 2

Contrarian Trading Strategy

For the daily returns analysis, we simulate a specific strategy – first proposed by Lehmann (1990) and Lo and MacKinlay (1990) – which we apply directly to US equities and subsets of US equities (e.g., stocks that trade primarily on the NYSE). Given a collection of M securities, consider a long/short equity market-neutral strategy consisting of an equal dollar amount of long and short positions. At each rebalancing interval, the long positions consist of "losers" (underperforming stocks, relative to the equal-weighted market average) and the short positions consist of "winners" (outperforming stocks, relative to the same market average). Specifically, if $w_i[n]$ is the portfolio weight of security i at date n, then,

$$w_i[n] = -\frac{1}{M}(r_i[n-k] - r_m[n-k]),\tag{2.1}$$

$$r_m[n-k] = \frac{1}{M}\sum_{i=1}^{M} r_i[n-k]\tag{2.2}$$

for some $k > 0$.

Note that the portfolio weights are the negative of the degree of outperformance k periods ago, so each value of k yields a somewhat different strategy. For our purposes, we set $k = 1$ period, where the length of each period will depend on the context (e.g., daily vs. intraday returns analysis). By buying the previous period's losers and selling the previous period's winners at the onset of each period, this strategy actively bets

on mean reversion across all M stocks and profits from reversals that occur within the subsequent interval. For this reason, this strategy has been called "contrarian" as it benefits from market overreaction and mean reversion (i.e., when underperformance is followed by positive returns and outperformance is followed by negative returns). By construction, the weights sum to zero and therefore the strategy is considered a "dollar-neutral" or "arbitrage" portfolio. Accordingly, the weights have no natural scale since any constant multiple of the weights will also sum to zero. In practice, however, the return of such a strategy over any finite interval is easily calculated as the profit-and-loss of that strategy's positions over the interval divided by the capital required to support those positions. Under Regulation T, the minimum amount of capital required is one-half of the total capital invested (often stated as 2:1 leverage, or a 50% margin requirement) and therefore the unleveraged (Reg T) portfolio return, $r_p[n]$ is given by:

$$r_p[n] = \frac{\sum_{i=1}^{M} w_i[n] r_i[n]}{I[n]} \tag{2.3}$$

$$I[n] = \frac{1}{2} \sum_{i=1}^{M} |w_i[n]|. \tag{2.4}$$

Since the portfolio weights are proportional to the differences between the equal-weighted market index and the returns, securities that deviate more positively from the market at time $n - k$ will have a greater negative weight in the time n portfolio, and vice-versa. This means the strategy bets on greater or more certain price reversals following large price changes. Such a strategy is designed to take advantage of negative serial autocorrelation, but Lo and MacKinlay (1990) show that, if returns are positively cross-autocorrelated, then a contrarian investment strategy can also yield positive profits.

For the intraday returns analysis, we apply a similar contrarian trading simulation to transactions data from November 1, 1994 to December 31, 2012 for the stocks in the S&P 1500 universe. This high-frequency mean reversion strategy has been simplified relative to the daily returns analysis to improve computational robustness and is

24

based on buying losers and selling winners over lagged m-second returns, where we vary m from 15 to 1800 seconds. Specifically, each trading day is broken into non-overlapping m-second intervals, and during each m-second interval we construct a long/short dollar-neutral portfolio that is long those 80 stocks with returns that rank lowest (dropping the 5 lowest) over the previous m-second interval, and short those 80 stocks with returns that rank highest (dropping the 5 highest) over the previous m-second interval. Stocks are equal-weighted, and no overnight positions are allowed. The value of the portfolio is then calculated for the next m-second holding period, and this procedure is repeated for each of the non-overlapping m-second intervals during the day. We always use the last traded price in each m-second interval to calculate returns, and so the first set of prices for each day are the prices based on trades just before 9:30am plus m seconds, and the first set of positions are established at 9:30am plus $2m$ seconds. We only use those transactions that occur during normal trading hours, and discard all trades that are reported late, out of sequence, or have a non-zero correction field (see section 1.3 for further details).

2.1 Expected Profits

Making minor modifications to the profitability analysis developed by Lo and MacKinlay (1990), we can estimate the fraction of the expected contrarian trading strategy profits due to the cross-covariance versus the autocovariance of returns. Rearranging Eq. (2.3), letting $\tilde{r}_i[n] = r_i[n]/I_k[n]$ and taking expectations yields the following:

$$E[r_p[n]] = \frac{\mathbf{1}^T \tilde{\Gamma}_k \mathbf{1}}{M^2} - \frac{1}{M} tr(\tilde{\Gamma}_k) - \frac{1}{M} \sum_{i=1}^{M} (\mu_i - \mu_m)(\tilde{\mu}_i - \tilde{\mu}_m) \tag{2.5}$$

where $\mu_i = E[r_i[n]]$, $\mu_m = E[r_m[n]]$, $\tilde{\mu}_i = E[\tilde{r}_i[n]]$, $\tilde{\mu}_m = E[\tilde{r}_m[n]]$ and $\tilde{\Gamma}_k(i,j) = E[r_i[n]\tilde{r}_j[n-k]] - E[r_i[n]]E[\tilde{r}_j[n-k]]$. Similar to the derivation in Lo and MacKinlay (1990), $E[r_p[n]]$ can be decomposed into three terms: one $(\tilde{C}_k)$ that depends only on cross-covariance terms, a second $(\tilde{O}_k)$ that depends only on autocovariance terms, and

25

a third $[\sigma^2(\tilde{\mu})]$ that is independent of the auto- and cross-covariances. Specifically,

$$E[r_p[n]] = \tilde{C}_k + \tilde{O}_k - \sigma^2(\tilde{\mu}) \tag{2.6}$$

where

$$\tilde{C}_k = \frac{1}{M^2}[\mathbf{1}^T\tilde{\Gamma}_k\mathbf{1}-tr(\tilde{\Gamma}_k)] \qquad \tilde{O}_k = -\frac{M-1}{M^2}tr(\tilde{\Gamma}_k) \qquad \sigma^2(\tilde{\mu}) = \frac{1}{M}\sum_{i=1}^{M}(\mu_i-\mu_m)(\tilde{\mu}_i-\tilde{\mu}_m).$$
$$\tag{2.7}$$

In Appendix A we show how to derive the sampling theory for the estimators $\hat{\tilde{C}}_k$, $\hat{\tilde{O}}_k$, and $\hat{\sigma}^2(\tilde{\mu})$.

2.2 A Model of Returns

As in Khandani and Lo (2011), suppose that stock returns satisfy the following simple linear multivariate factor model:

$$r_i[n] = \mu_i + \beta_i\nu[n] + \lambda_i[n] + \eta_i[n], \tag{2.8}$$

$$\lambda_i[n] = \theta_i\lambda_i[n-1] - \varepsilon_i[n] + \varepsilon_i[n-1], \qquad \theta_i \in (0,1), \tag{2.9}$$

$$\nu[n] = \rho\nu[n-1] + \zeta[n], \qquad \rho \in (-1,1), \tag{2.10}$$

where $\eta_i[n]$, $\varepsilon_i[n]$, and $\zeta[n]$ are white-noise random variables that are uncorrelated at all leads and lags.

The impact of random idiosyncratic supply and demand shocks is represented by $\varepsilon_i[n]$, the variance of which, $\sigma_{\varepsilon_i}^2$, is representative of the width of the bid-ask spread. Notice that Eq. (2.9) is an ARMA(1,1) process and exhibits negative autocorrelation to capture the mean reversion generated by market making activity such as the bid-ask bounce, as in Roll (1984), and market maker inventory rebalancing, as in Hasbrouck and Sofianos (1993). Additionally, decreasing the parameter θ_i increases the speed of mean reversion and is representative of less persistence in buying/selling pressures.

Mean reversion or momentum in the market common factor, $\nu[n]$, is captured

by the parameter ρ when it is less than or greater than zero, respectively. Note that this model does not consider idiosyncratic autocorrelation induced from stock-specific fluctuations, $\eta_i[n]$, which are unrelated to liquidity or common factors (e.g., stock specific investor overreaction or under reaction followed by a correction).

By changing the variances of our white noise random variables, we can change the relative importance of market making activity to informational shocks with respect to price movements in stock i. For example, increasing the power in $\varepsilon_i[n]$ could be representative of wider spreads and decreased liquidity, while increasing the variance of $\zeta[n]$ would increase the importance of common factor turbulence. Changing the variance of $\eta_i[n]$ would not change the expected profits of the contrarian trading strategy since it is not cross or serially correlated with returns at any leads or lags, but would substantially increase the volatility of the strategy's returns if M, the number of stocks, were small. We shall see that both cases, where liquidity and common factors dominate, significantly affect contrarian trading profitability in our daily and intraday returns analysis.

Lo and Khadani (2011) show that the expected one-period profit of Eq. (2.3) with $k = 1$ and without the normalization of $I[n]$ is given by,

$$E[r_p[n]] = \frac{M-1}{M^2} \sum_{i=1}^{M} \frac{1-\theta_i}{2} \sigma_{\lambda_i}^2 - \rho \sigma_\nu^2 \sigma^2(\boldsymbol{\beta}) - \sigma^2(\boldsymbol{\mu}), \qquad (2.11)$$

where $\boldsymbol{\beta}$ and $\boldsymbol{\mu}$ are (M × 1)-vectors representing the collection of β_i's and μ_i's, respectively.

When $\rho > 0$ and there is dispersion in $\boldsymbol{\beta}$, then the momentum in the common factor causes a decrease in contrarian trading profits. In our subsequent analysis, we shall see that the common factor contribution is usually dominated by the market making component, in part because market neutrality ensures that profits remain largely insulated from common factor effects. In certain cases, when market pressures build (e.g., in the event of a large factor unwind), then this term can overwhelm the liquidity term and drive profits negative. If however $\rho < 0$, then the negative autocorrelation in the common factor contributes positively to expected profits. Assuming

$\sigma^2(\boldsymbol{\mu}) \approx 0$, then the expected profits will be unambiguously positive as the strategy always benefits from the mean reversion induced by market making activity.

Chapter 3

Daily Returns Analysis

In this section, we summarize the contrarian trading strategy's historical performance over the past five decades as an interesting and informative case study of how market making activity can alter price dynamics. Panel A of Fig. 3-1 illustrates the average daily returns of the contrarian strategy when applied to stocks sorted into NYSE and NASDAQ NMS portfolios by their CRSP exchange code identifier. Between 1983 and 1988, the average daily returns for the NASDAQ NMS portfolio increased from -0.67% to 1.64% while the NYSE portfolio remained relatively unchanged and unprofitable. Throughout the late-1980s and early-1990s the profitability of the contrarian strategy remained just below 2% for NMS stocks, until 1995 when it began to gradually decay towards 0% over the next decade. During the 2008 financial crisis and the ensuing downturn, profitability for the strategy again increased to about 0.7% for NMS stocks but remained relatively flat for NYSE stocks.

Panel B of Fig. 3-1 decomposes the average contrarian trading profits of the NMS portfolio into its autocorrelation and cross-correlation components as explained in Eq. (2.6). We notice the autocorrelation term is almost exclusively responsible for changes in profitability over time which suggests that profits are driven either by market making activity (e.g., a larger bid-ask bounce due to wider spreads) or by common factor negative serial autocorrelation. In Panel C, to test the common factor hypothesis, we plot both the one-day autocorrelation of equal-weighted NMS returns (ρ_{mkt}) which can be considered a proxy to a market information common factor as well as the

cross-sectional mean of individual security return one-day autocorrelations $(\bar{\rho}_i)$. We see that initially, contrarian profits are negatively correlated with the idiosyncratic autocorrelation, and less closely related to the market autocorrelation, suggesting that the liquidity provision term in Eq. (2.11) is dominant and driving profitability. During the 2008 financial crisis and thereafter, contrarian profits appear to be more closely related to market autocorrelation suggesting that the common factor term in Eq. (2.11) began to play a relatively more substantial role in determining expected profits.

Tables 3.1 and 3.2 show that the composition of the NMS portfolio, in terms of number of stocks and average market capitalization, rapidly changes over its first couple of years in existence. Initially, only the most active NASDAQ stocks had their recent trade and volume data relayed electronically on the NMS. Between 1983 and 1985 the composition of the NMS changed as less actively traded stocks were added, causing the contrarian profitability during this period to rise. By 1986 the number of stocks as well as the average market capitalization had stabilized, suggesting that the subsequent rise and fall in contrarian profitability was related to other factors. Table 3.1 also reports the strategy's annualized Sharpe ratio relative to a 0% risk-free rate as a simple measure of the strategy's expected return per unit of risk. As mentioned in Khandani and Lo (2007), some of these Sharpe ratios may seem unrealistically high, but it should be kept in mind that we have not incorporated the formidable transaction costs required to rebalance each stock daily. However, if these returns are capturing the bid-ask spread through liquidity provision, it is reasonable that a dampened form of these returns could have been attained by market making statistical arbitrage strategies given the economics of price discovery. Finally, we note that the increase in average daily returns for the NMS portfolio around 2008 occurred simultaneously with an increase in the strategy's volatility. The earlier period of high profitability during the late 1980s and early 1990s did not have this feature, suggesting alternative mechanisms were driving returns in the two periods.

30
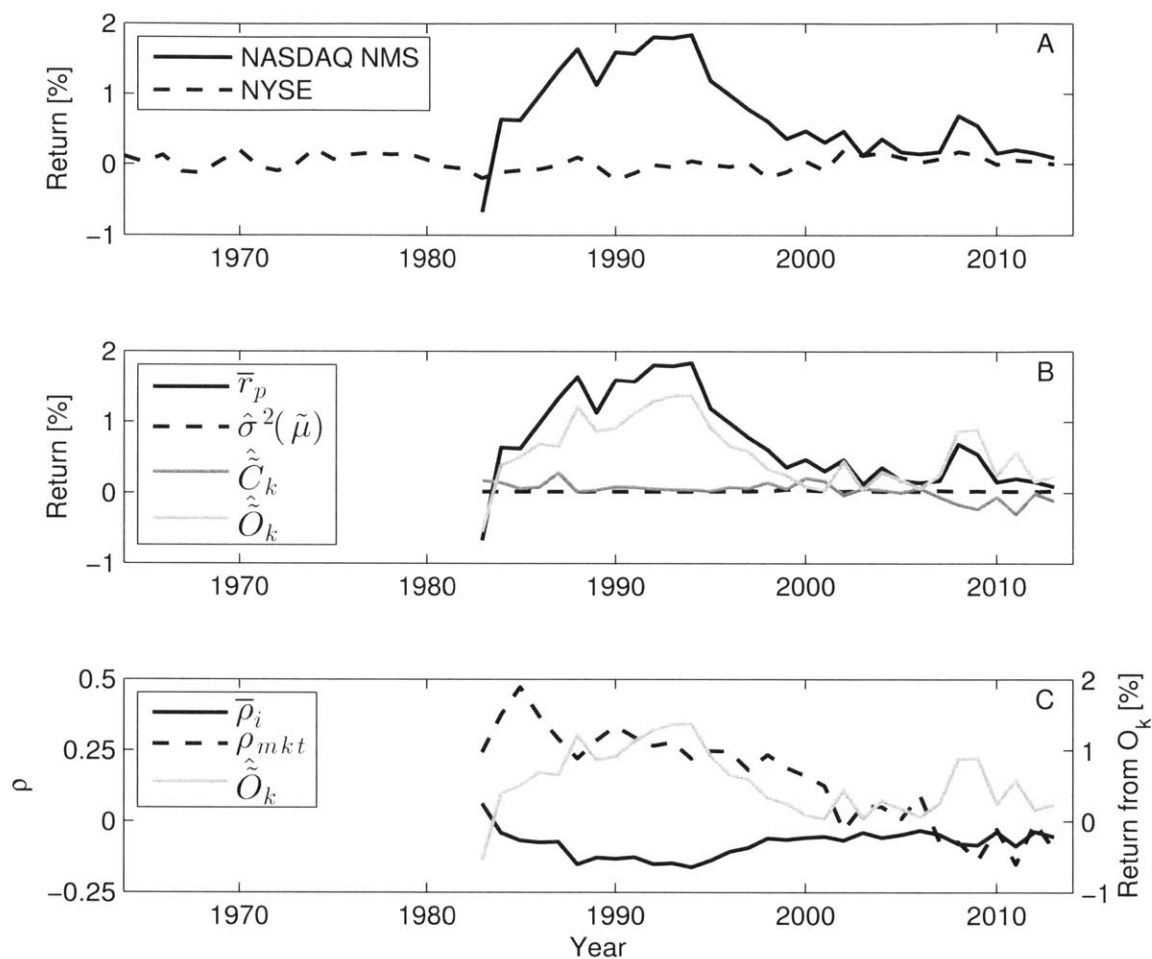
Figure 3-1: Year-by-year average daily returns of Lo and MacKinlay's (1990) contrarian trading strategy applied to all US common stocks (CRSP share codes 10 and 11) and exchange-sorted subsets with share prices above \$5 and less than \$2,000 from January 1, 1964 to December 31, 2013 (Panel A). Decomposition of contrarian NMS portfolio profits into autocovariance and cross-covariance components (Panel B). Idiosyncratic and market autocorrelation estimates over the same time interval (Panel C).

Table 3.1: Year-by-year average daily returns, standard deviations of daily returns, and annualized Sharpe ratios ($\sqrt{250}\times$ average daily return/standard deviation) of Lo and MacKinlay's (1990) contrarian trading strategy applied to all US common stocks (CRSP share codes 10 and 11) with share prices above \$5 and less than \$2,000 within exchange-sorted portfolios from January 1, 1964 to December 31, 2013.

Year	Avg daily returns (%) NYSE	Avg daily returns (%) NMS	Std daily returns (%) NYSE	Std daily returns (%) NMS
1964	-0.12	-	0.40	-
1965	0.03	-	0.67	-
1966	0.13	-	0.77	-
1967	-0.10	-	0.53	-
1968	-0.12	-	0.49	-
1969	0.05	-	0.49	-
1970	0.20	-	0.60	-
1971	-0.03	-	0.49	-
1972	-0.09	-	0.39	-
1973	0.00	-	0.64	-
1974	0.22	-	0.60	-
1975	0.07	-	0.53	-
1976	0.14	-	0.40	-
1977	0.16	-	0.32	-
1978	0.14	-	0.65	-
1979	0.15	-	0.50	-
1980	0.07	-	0.64	-
1981	-0.03	-	0.47	-
1982	-0.06	-	0.55	-
1983	-0.19	-0.67	0.46	1.41
1984	-0.11	0.64	0.42	0.69
1985	-0.08	0.63	0.38	0.54
1986	-0.07	0.98	0.50	0.59
1987	-0.01	1.33	0.89	1.11
1988	0.10	1.64	0.48	0.51

Year	Avg daily returns (%) NYSE	Avg daily returns (%) NMS	Std daily returns (%) NYSE	Std daily returns (%) NMS
1989	-0.02	1.13	0.54	0.47
1990	-0.22	1.59	0.55	0.66
1991	-0.12	1.57	0.48	0.62
1992	-0.01	1.81	0.44	0.69
1993	-0.04	1.79	0.45	0.46
1994	0.05	1.84	0.38	0.47
1995	-0.01	1.19	0.42	0.62
1996	-0.03	0.99	0.45	0.59
1997	0.02	0.78	0.60	0.79
1998	-0.19	0.61	0.74	1.07
1999	-0.11	0.37	0.80	1.18
2000	0.04	0.47	1.03	1.85
2001	-0.09	0.31	0.96	1.65
2002	0.20	0.47	0.79	1.14
2003	0.13	0.12	0.53	0.64
2004	0.17	0.36	0.49	0.61
2005	0.10	0.18	0.50	0.49
2006	0.03	0.15	0.54	0.54
2007	0.07	0.18	0.79	0.72
2008	0.18	0.69	1.88	1.42
2009	0.13	0.55	1.56	1.16
2010	-0.00	0.16	0.66	0.63
2011	0.06	0.20	0.82	0.85
2012	0.04	0.16	0.57	0.53
2013	0.00	0.09	0.47	0.51

Table 3.1: (Continued)

Year	Annualized Sharpe ratio NYSE	Annualized Sharpe ratio NMS	Avg number of stocks NYSE	Avg number of stocks NMS
1964	4.71	-	1042	0
1965	0.71	-	1083	0
1966	2.68	-	1104	0
1967	-3.06	-	1150	0
1968	-3.95	-	1153	0
1969	1.55	-	1157	0
1970	5.18	-	1143	0
1971	-1.07	-	1214	0
1972	-3.84	-	1264	0
1973	0.03	-	1246	0
1974	5.84	-	1127	0
1975	2.11	-	1141	0
1976	5.45	-	1239	0
1977	8.01	-	1309	0
1978	3.41	-	1321	0
1979	4.70	-	1328	0
1980	1.64	-	1325	0
1981	-1.12	-	1326	0
1982	-1.80	-	1271	0
1983	-6.61	-7.53	1311	193
1984	-4.03	14.59	1289	632
1985	-3.51	18.31	1260	1005
1986	-2.13	26.09	1235	1418
1987	-0.24	18.92	1213	1503
1988	3.32	51.23	1154	1313
1989	-0.61	38.35	1130	1268
1990	-6.31	38.16	1084	1119
1991	-3.92	40.01	1099	1082
1992	-0.34	41.13	1217	1321
1993	-1.47	61.14	1340	1719
1994	2.06	61.38	1443	1945
1995	-0.22	30.18	1483	1863
1996	-1.21	26.30	1566	2280
1997	0.43	15.54	1672	2529
1998	-4.00	9.02	1701	2416
1999	-2.11	4.88	1629	2201
2000	0.62	4.06	1460	2201
2001	-1.42	2.98	1347	1735
2002	4.03	6.54	1322	1530
2003	3.83	2.99	1317	1500
2004	5.37	9.30	1380	1755
2005	3.10	5.68	1387	1742
2006	0.78	4.27	1374	1786
2007	1.49	3.87	1357	1815
2008	1.50	7.70	1239	1494
2009	1.31	7.43	1090	1201
2010	-0.03	3.96	1189	1345
2011	1.11	3.81	1206	1343
2012	1.03	4.78	1187	1258
2013	0.02	2.82	1215	1311

Table 3.2: Year-by-year average market capitalizations and share prices of US common stocks (CRSP share codes 10 and 11) with share prices above $5 and below $2,000 within exchange-sorted portfolios from January 1, 1964 to December 31, 2013.

Year	Avg market cap ($MM)		Avg price ($)	
	NYSE	NMS	NYSE	NMS
1964	408	-	38.11	-
1965	436	-	38.45	-
1966	426	-	36.70	-
1967	459	-	39.94	-
1968	504	-	41.82	-
1969	517	-	36.51	-
1970	450	-	27.30	-
1971	524	-	31.94	-
1972	579	-	32.75	-
1973	577	-	27.16	-
1974	486	-	21.09	-
1975	519	-	22.49	-
1976	584	-	24.81	-
1977	561	-	23.72	-
1978	566	-	24.03	-
1979	618	-	24.73	-
1980	729	-	26.88	-
1981	813	-	27.11	-
1982	796	-	24.28	-
1983	1059	445	30.89	25.45
1984	1056	181	27.13	16.83
1985	1248	171	29.70	18.76
1986	1570	188	33.37	20.28
1987	1868	204	32.40	17.85
1988	1786	206	28.22	16.29

Year	Avg market cap ($MM)		Avg price ($)	
	NYSE	NMS	NYSE	NMS
1989	2122	233	30.83	17.51
1990	2230	242	28.16	16.12
1991	2493	309	29.91	18.61
1992	2597	326	29.71	18.50
1993	2681	329	30.28	18.53
1994	2607	333	28.20	17.64
1995	2958	443	29.13	19.44
1996	3519	487	31.78	19.48
1997	4253	597	35.18	20.38
1998	5135	804	34.19	20.38
1999	6132	1357	30.37	22.07
2000	7139	2191	30.80	25.92
2001	7308	1566	30.08	19.77
2002	6535	1318	28.87	19.28
2003	6373	1439	28.65	20.78
2004	7271	1535	33.27	21.63
2005	7911	1644	35.94	22.35
2006	8640	1725	37.57	23.35
2007	9628	1901	40.38	23.82
2008	8298	1846	33.56	20.79
2009	7028	1991	28.42	19.56
2010	7863	2316	33.63	22.51
2011	8734	2651	37.46	25.13
2012	9173	3205	38.94	26.91
2013	10679	3677	46.28	31.77

3.1 Bid-Ask Spreads and Even-Eighth Quotes

To determine what features of the market microstructure were driving changes in liquidity provision profitability, we analyze the characteristics of the inside spread and quotes. Panel A of Fig. 3-2 applies the contrarian strategy to three market capitalization tercile subsets of the NMS portfolio and Panel C shows the cross-sectional average of their percentage bid-ask spreads. The small cap stocks generally exhibit larger spreads relative to their prices, and this spread is translated into larger contrarian strategy profits. From these two panels, it appears the dynamics causing negative autocovariance in the returns of NMS stocks are directly related to the width of the percentage bid-ask spread. By implementing the same contrarian strategy on individual stock returns that have been calculated using inside bid-ask averages instead of transaction prices we see, in Panel B, that the profitability for all subsets substantially decreases. Finally, in absolute terms, Panel D shows that the average bid-ask spread is in fact similar across all three market cap terciles except during the financial crisis when increased market making costs induced wider spreads, especially for the thinly traded small cap stocks. Unlike the financial crisis, the wide spreads seen in the early 1990s did not occur during a period of historically high market volatility leading us to ask: what drove the widening of spreads, and consequently the increased contrarian trading strategy profits, post-1987, and moreover, why did they decrease in the mid-to-late 1990s?

To shed light on these questions we turn to Christie and Schultz (1994) and Christie, Harris and Schultz (1994) who provide evidence that market makers of active NASDAQ stocks were implicitly colluding to maintain wide spreads by avoiding odd-eighth quotes. Fig. 3-3 compares the fraction of even-eighth and odd-sixteenth inside bid-ask quotes and transaction prices for our market-cap tercile NYSE and NMS portfolios. To prevent clutter, we leave out the fraction of odd-eighth quotes as it can be calculated before 2001 as one minus the fraction of even-eighth and odd-sixteenth quotes[1]. Since the characteristics of both bid and ask quotes are nearly identical,

---

[1] In 2001, both NYSE and NASDAQ switched to decimalization making odd-eighth and odd-sixteenth quotes unattainable. Also note that both exchanges changed their minimum price varia-
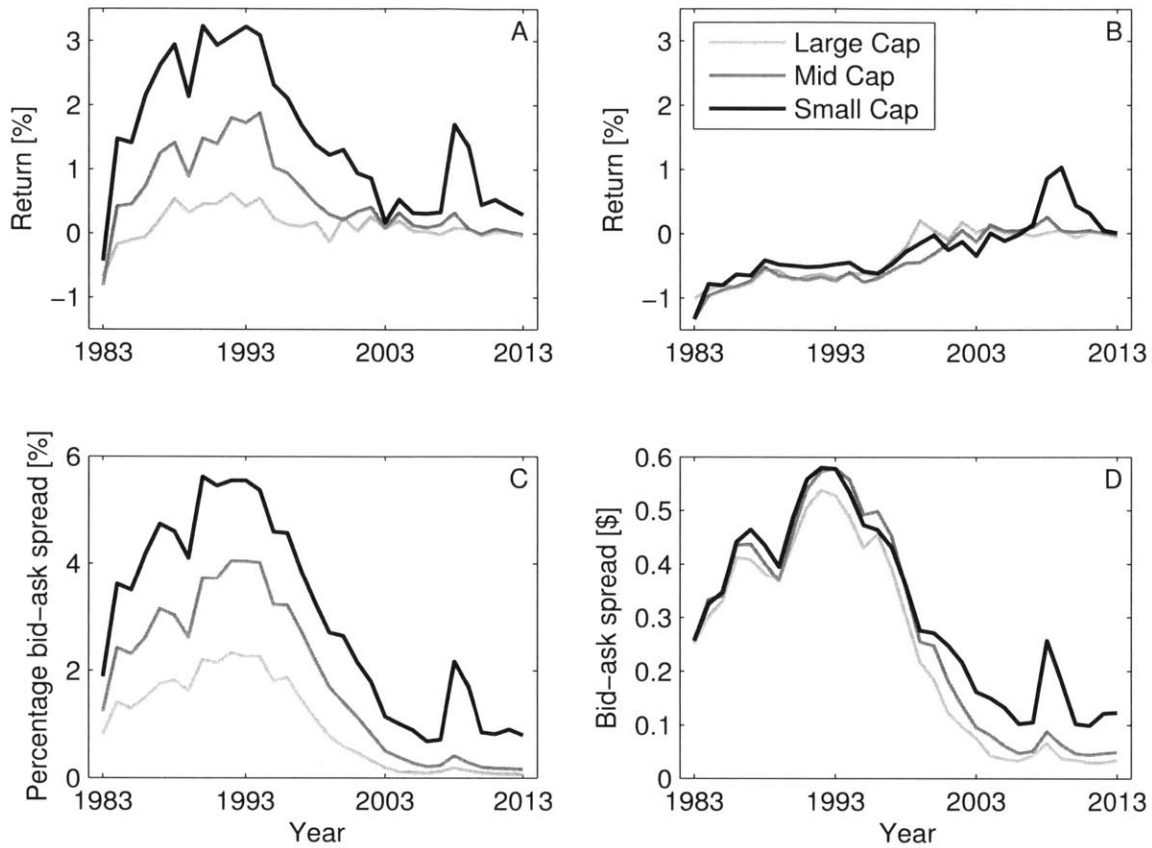
Figure 3-2: Year-by-year average daily returns of the contrarian trading strategy applied to US common NASDAQ NMS market capitalization terciles with share prices above $5 and less than $2,000 from January 1, 1983 to December 31, 2013 (Panel A). Same analysis as in Panel A except using the mean of the inside bid and ask quotes instead of transaction prices to calculate daily stock returns (Panel B). Cross-sectional average of percentage bid-ask spread (Panel C) and bid-ask spread (Panel D) over the same time interval.

they have been plotted together. For the small cap NMS portfolio, the fraction of even-eighth quotes initially rises from 0.65 to 0.74, and may be explained by the rapidly changing composition of the NMS as less actively traded stocks were added to the system. Between 1989 and 1993, a period in which the number of stocks in the NMS had become relatively stable, the fraction of even-eighth quotes increased further from 0.75 to 0.85! Post-1994, perhaps due in part to the publication of the Christie and Schultz findings, the fraction of even-eighth quotes started to decline but remained well above the value of 0.5 that would be expected in an efficient market

tions to from eighths to sixteenths of a dollar for all stocks in 1997.

where the distribution across even and odd quotes is uniform. Once sixteenths were introduced in 1997, the fraction even-eighth eighth quotes continued to decline to about 0.5 but remained well above the new efficient market value of 0.25. Post-2001, after decimalization had been introduced, the fraction of even-eighth quotes dropped rapidly towards its efficient market value of 0.04 but even then there was some latency in the decline. In fact, for the small cap NMS stocks there was a slight preference for even-eighth quotes during the 2008 financial crisis. The same general pattern can be seen across each market cap-based NMS portfolio and is consistent with the result shown in Fig. 3-2 where average bid-ask spreads, in absolute terms, were approximately equal across each tercile. Panel B shows that the pattern is roughly the same for transaction prices with slightly smaller values for the fraction of even-eighth prices during the initial period.

Panels B and D show the case for the NYSE cap-based terciles. A continuous series of bid and ask quote data are not available on CRSP for these stocks until December 28, 1992, and so we focus our attention on the transaction price data noting that, where the bid and ask quote data are available, the results are similar. Starting in 1982, the fraction of even-eighth prices is 0.58, which is already much closer the expected efficient market value of 0.5. Over the next 14 years the fraction of even-eighth prices declines gradually and almost linearly to 0.53 in 1996. In 1997, the NYSE changed its minimum price variation from eighths to sixteenths of a dollar and almost immediately by 1998 the fraction of even-eighths decreased to 0.33 and the fraction of odd-sixteenths increased to about 0.4, close to their efficient market values of 0.25 and 0.5, respectively. This is in sharp contrast to the pattern of slow adoption of odd-sixteenths seen in the NMS subsets. Post-2001, once decimalization was introduced, the fraction of even-eighths again dropped sharply toward its efficient value of 0.04, but similar to the NMS portfolios, there was some latency in this decline.

A higher fraction of even-eighths in both quotes and transaction prices, delays in integrating new price levels, as well as a larger profitability of the liquidity providing contrarian strategy all provide evidence of inefficient market making dynamics in the NMS relative to the NYSE. Competition for order flow amongst market makers was
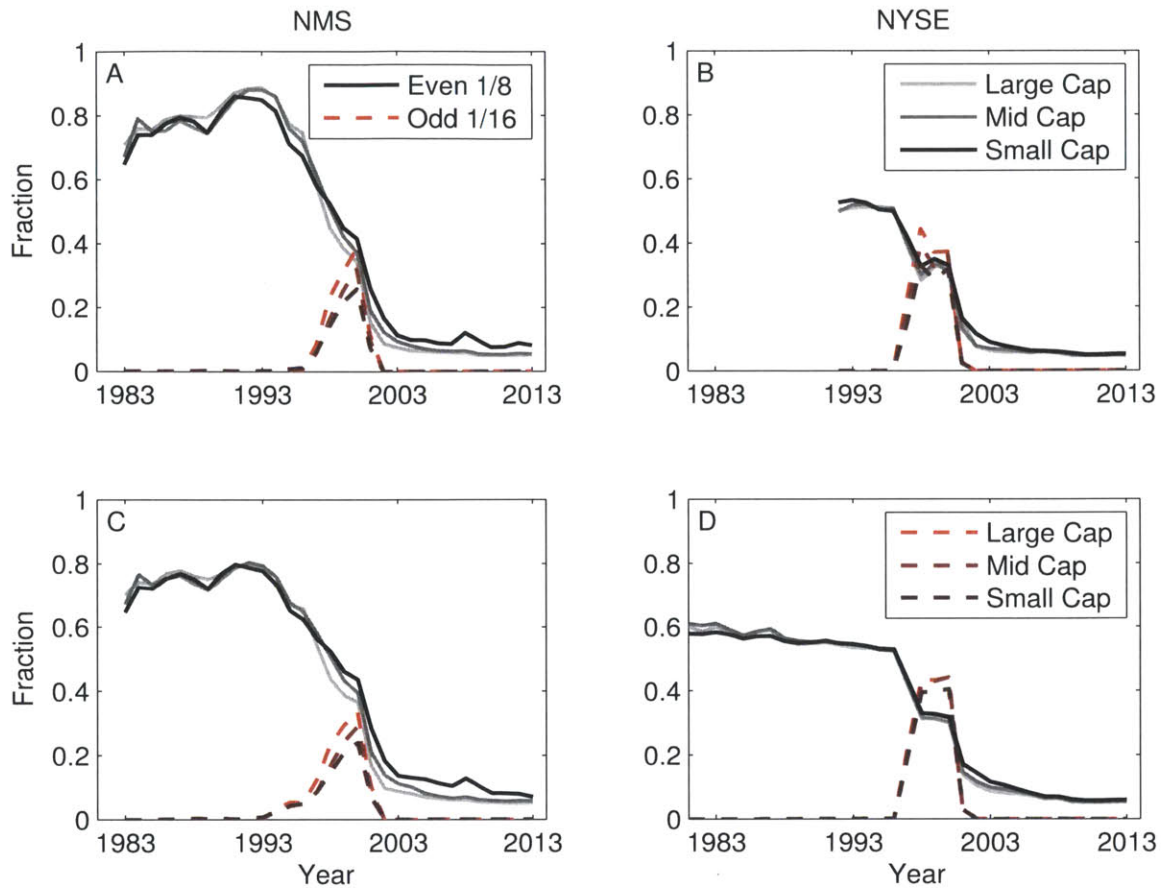
Figure 3-3: Year-by-year cross-sectional average of the fraction of even-eighth inside quotes within US common NASDAQ NMS and NYSE market capitalization terciles with share prices above $5 and less than $2,000 from January 1, 1983 to December 31, 2013 (Panels A and B respectively). Same analysis as in Panels A and B except applied to transaction prices instead of inside quotes (Panels C and D, respectively).

designed to encourage tighter spreads, however a seemingly implicit agreement to offer only even-eighth quotes would have maintained the bid-ask spread to a minimum of $0.25, driving up profits for liquidity providers. Perhaps part of the reason for the widening spreads was due to increased market making risk associated with the growing presence of institutional trading as described in Kothare and Laux (1995). Or maybe institutional traders found trading in even-eighths more efficient and were directly responsible for the rise in even-eighth quotes and transaction prices. It does appear though, that the avoidance of odd-eighth quotes was in place even before institutional trading was a significant fraction of the volume on the NMS. Moreover,

abuse of the retail trading system SOES is also unlikely to explain the full extent of these widening spreads seen across all market cap cross-sections as the professional day traders limited their trading activity to the most active NMS stocks.

One question that, due to limitations in our dataset, we cannot conclusively answer is whether the avoidance of odd-eighth fractions was facilitated by the real-time reporting procedures of the NMS or if it was prevalent amongst all NASDAQ securities. Prices for non-NMS NASDAQ stocks prior to June 15, 1992 are recorded in our dataset as the mean of the inside bid-ask quotes and the quotes themselves are not recorded. Post June 15, 1992 these securities adopted a similar real-time transaction reporting scheme to the NMS securities making a direct comparison uninteresting. For both NMS and non-NMS NASDAQ securities, we looked at the average of inside bid and ask quotes noting that a greater fraction of odd-eighth quotes would translate into a greater prevalence of odd-sixteenth midpoints. The similarity between the two exchanges, as seen in Fig. 3-4, and the overall scarcity of odd-sixteenth midpoints provides partial evidence that the tendency to avoid odd-eighth quotes was not limited to the NMS. Note that the rapidly growing and changing composition of the NMS through the inclusion of less actively traded stocks is likely the cause of the initial decrease in the fraction of odd-sixteenth midpoints in Panel A.

In 1994, the tendency for NMS market makers to avoid odd-eighth quotes, the width of the average inside spread and the profitability of the contrarian strategy all started to decline. A breakdown of implicit pricing agreements amongst market makers may well have been triggered by the negative publicity surrounding the Christie and Schultz (1994) findings. Indeed, the follow up paper by Christie, Harris and Schultz (1994) reported that the average inside spread for Amgen, Cisco Systems, and Microsoft fell to about $0.15 per share after varying between $0.25 and $0.45 over the preceding 17 months. Competition from non-traditional market makers also played a role in narrowing spreads across all market cap-based NMS portfolios as hedge funds and proprietary trading desks, attracted by wide spreads and the efficiency afforded by electronic trading networks, injected enormous amounts of liquidity into the market. These new entrants would have traded within the even-quoted spread
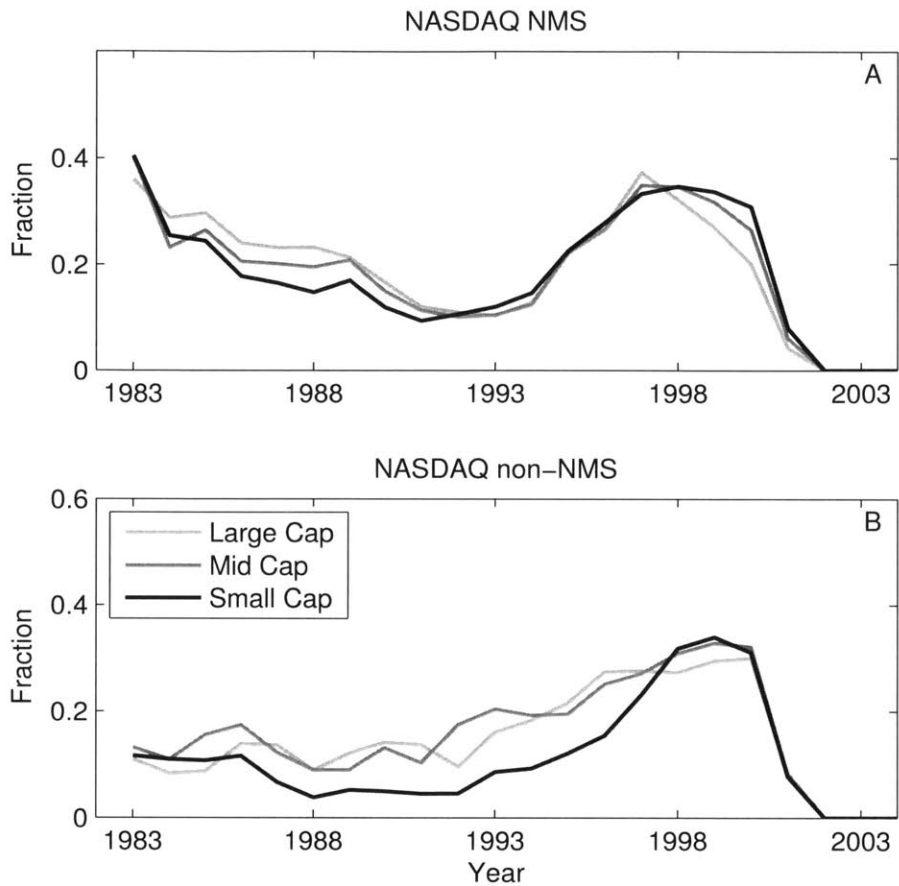
Figure 3-4: Year-by-year cross-sectional average of the fraction of odd-sixteenth inside bid-ask quote midpoints within US common NASDAQ NMS (Panel A) and NASDAQ non-NMS (Panel B) market capitalization terciles with share prices above $5 and less than $2,000 from January 1, 1983 to December 31, 2003.

in an effort to attract more order flow leading to the eventual collapse of any implicit pricing agreements that may have existed amongst traditional market makers. Finally, we highlight the small resurgence in the preference for even-eighth quotes for small cap NMS stocks during the 2008 financial crisis. This observation suggests that market makers may default to a limited subset of quotes to widen spreads in response to thinning liquidity and increased inventory and asymmetric information risks.

Chapter 4

Intraday Returns Analysis

Having documented the historical behavior of market making profits for daily returns, we now turn our attention to intraday returns by applying a similar analysis to the transactions dataset from November 1, 1994 (the first month after the launch date of the S&P 600) to December, 31, 2012 for stocks in the S&P 1500 universe.

Fig. 4-1 plots the 5-day moving average of the daily returns of the simplified mean reversion strategy described in chapter 2 with non-overlapping, 600-second (10 minute) intervals for both NYSE and NASDAQ subsets of stocks. Dates that correspond to important shifts in these returns are also labeled and correspond to the following events:

1. June 2, 1997: NASDAQ changes its minimum tick size from 1/8th to 1/16th of a dollar.

2. June 24, 1997: NYSE changes its minimum tick size from 1/8th to 1/16th of a dollar.

3. January 29, 2001: NYSE changes its minimum tick size from 1/16th to 1/100th of a dollar.

4. April 9, 2001: NASDAQ changes its minimum tick size from 1/16th to 1/100th of a dollar.

5. September 11, 2001: September 11 terrorist attacks.

6. July 19, 2002: Stock-market downturn of 2002 related to the information technology bubble.

7. August 6, 2007: Quant Meltdown.

8. March 14, 2008: Bear Stearns suffers one-day loss of -47.4%.

9. October 6, 2008: Start of the worst week for the stock market in 75 years.

10. May 6, 2010: Flash Crash.

11. August 8, 2011: Stock markets fall due to fears of contagion of European sovereign debt crisis.
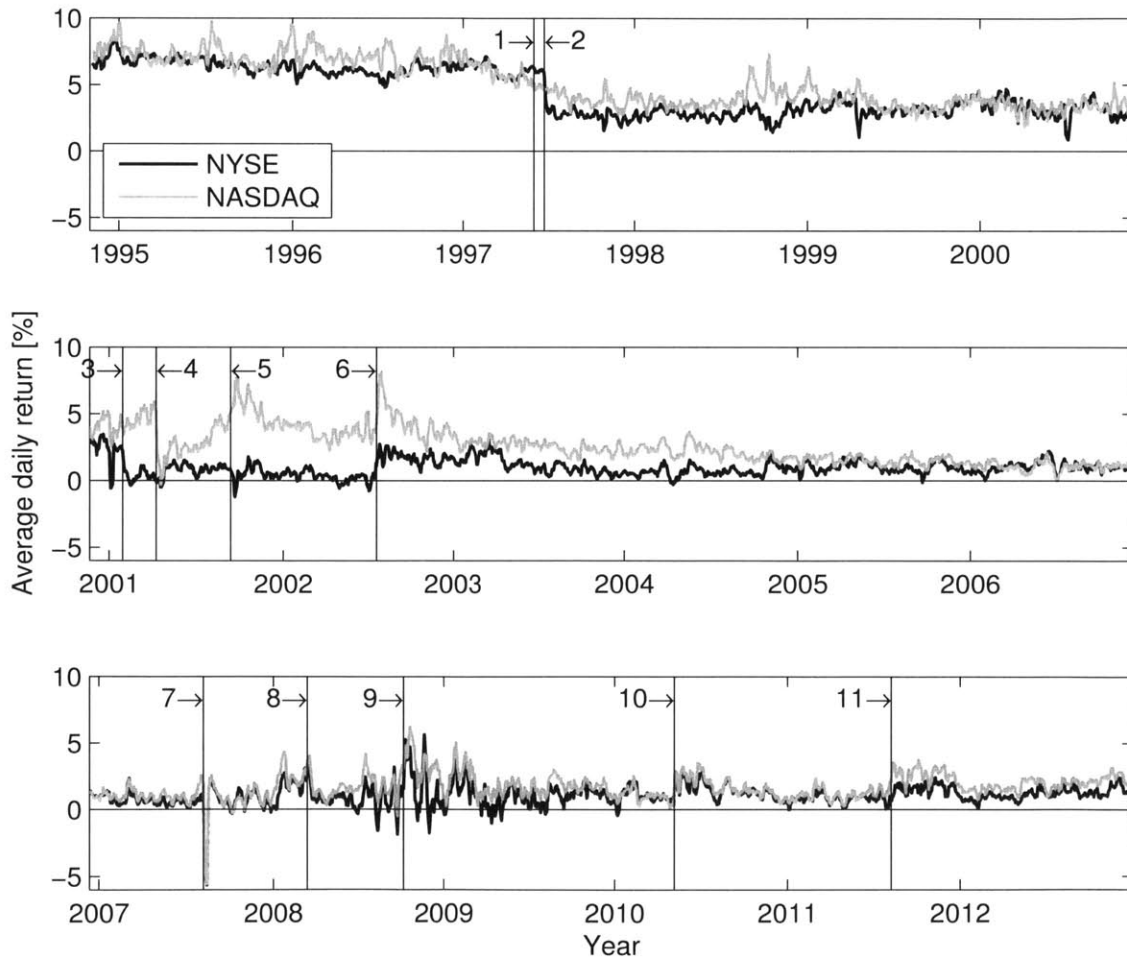

Figure 4-1: 5-day moving average of the daily returns of the mean reversion strategy with non-overlapping, 600-second intervals applied to exchange-sorted portfolios from November 1, 1994 to December 31, 2012.

42

Events 1 to 4 all correspond to market environment changes, where the minimum price variation (i.e., tick size) decreased. After each of these events, the average daily contrarian returns decreased indicating there had been a substantial reduction in the idiosyncratic mean reversion effect most likely due to a narrowing of spreads and a subsequently smaller bid-ask bounce. Note that, as documented in the daily returns analysis, the NASDAQ stocks show a gradual decline in contrarian trading profitability after their switch from eighths to sixteenths. As we discussed earlier, a strong bias towards even-eighth quotes hindered the adoption of odd-sixteenths in NASDAQ markets and led to a more gradual decrease in inside spreads. Over the next couple of years, this inefficiency appears to have been corrected as both NYSE and NASDAQ stocks show rapid decreases in profitability following their switch from sixteenths to decimalization in January and April 2001, respectively.

Throughout the 1990s and early 2000s, the profitability of the contrarian strategy applied to both NYSE and NASDAQ subsets consistently declined with the exception of a pair of increases in profitability caused by widening spreads following the September 11 terrorist attacks and the stock-market downturn of 2002. In fact, the fastest prolonged decay in profitability occurred between 1997 and 1999, a time when ECNs and algorithmic trading were rapidly growing in popularity. As we mentioned in the previous section, these advances enabled non-traditional market makers in the form of hedge funds and proprietary trading desks to compete with traditional market makers, adding enormous amounts of liquidity and narrowing the width of the bid-ask spread.

The monotonous downward trend in profitability during the mid-2000s was abruptly interrupted by the Quant Meltdown of 2007. This substantial deviation from the historical trend was marked by a sudden period of large losses followed by an equally rapid stabilization. The unprecedented nature of the losses helps to explain why so many previously consistently profitable quant funds were caught so severely exposed. As David Viniar, Chief Financial Officer of Goldman Sachs said: "We were seeing things that were 25-standard deviation moves, several days in a row" (Thal Larsen, 2007). Khandani and Lo (2007, 2011) provide evidence that this steep decline in

contrarian profits was triggered by the deleveraging of long/short equity positions ultimately causing the momentum in the common factor term in Eq. (2.11) to overwhelm the historically profit-generating market making activity term.

From 2008 to 2010, the pattern of profitability is highly variable and, on average, higher than in the previous period. The uncertainty surrounding financial markets in the wake of the housing bubble and the ensuing credit crisis increased volatility, reduced liquidity and widened spreads. In fact, as we saw in Fig. 3-3, some market makers reverted to their preference of even-eighth quotes as a means to widen their spreads, however, unlike their implicit price agreements in the early 1990s, this was in response to increased inventory and asymmetric information costs. As markets eventually stabilized, so did the profitability of the contrarian strategy until the Flash Crash of 2010 again caused liquidity to be withdrawn from the market. Healing occurred over the subsequent year until uncertainty surrounding the European sovereign debt crisis again drove contrarian profits up to levels where they have more or less remained the past couple of years.

To obtain a better sense of the performance of the mean reversion strategy across different time horizons, Fig. 4-2 plots the 250 trading-day moving average of normalized profits, risk and profits per unit risk for various values of m. In Panel A, the leverage of each strategy is altered such that they are all initially equally profitable which enables us to clearly visualize the evolution of their relative performance. Panel A demonstrates that the profitability of the longer, $m = 30$ min, time horizon strategy declines at a faster rate relative to the $m = 5$ and 10 min time horizon strategies between 1995 and 2003. Over the subsequent half decade the relative rates of decay are reversed as each strategy converges towards 0. This is consistent with the idea that non-traditional market makers were initially able to profit at longer time horizons and, as profit margins at these frequencies narrowed and technology improved, sought out profits in higher frequency domains. Panel B normalizes the risk, measured as the sampled standard deviation, such that one unit of risk is equivalent to the initial standard deviation of the $m = 30$ min trading strategy leveraged according to Panel A. Panel C plots the profitability from Panel A divided by the risk

44
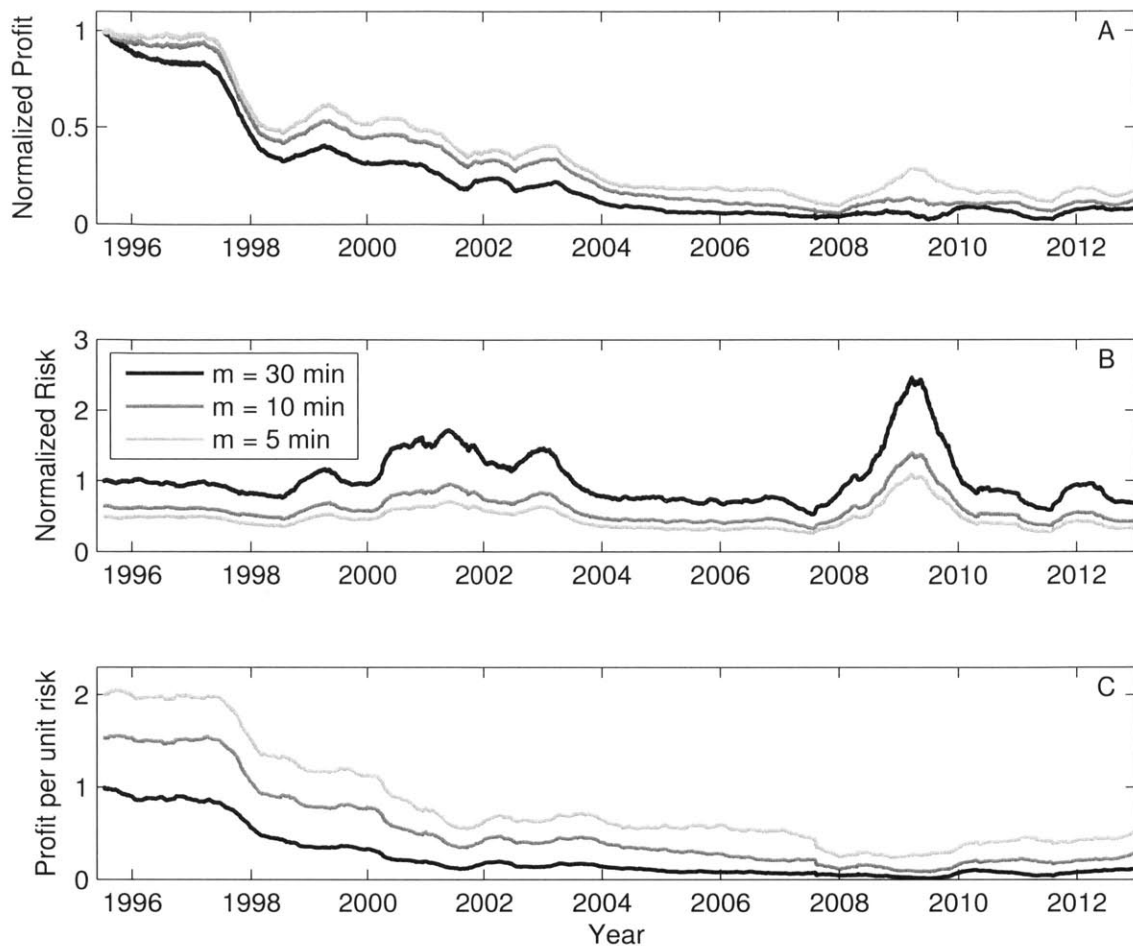
Figure 4-2: 250 trading day moving average of the contrarian strategy's daily returns for various time horizons (m) from November 1, 1994 to December 31, 2012 normalized so that each strategy's profits are initially equal (Panel A). The risk of each strategy described in Panel A is normalized such that the $m = 30$ min strategy initially holds one unit of risk (Panel B). The profit per unit risk for each strategy (Panel C).

in Panel B and shows that the longer time horizon strategies could only achieve the same average profitability as the shorter time horizon strategies by accepting higher levels of risk. The risk gap between the shorter and longer time horizon strategies in Panel B remained relatively constant throughout the entire sample period except during intervals marked by increased market instability. Turbulence and uncertainty surrounding the September 11, 2001 terrorist attack, 2002 stock market downturn and 2008 financial crisis impacted the volatility of the lower frequency strategies to a greater degree than they impacted the higher frequency strategies. This insulation

from market instability combined with the opportunity for substantially larger profits provided strong incentives for traders to shorten the time horizon of their strategies, ultimately leading to the development of today's HFT industry.

4.1 Analysis of the May 6, 2010 Flash Crash

In this section we increase the magnification of our analytical lens by another order of magnitude, and focus on the events of May 6, 2010, and compare them to those of August 2007. Panel A of Fig. 4-3 provides the minute-by-minute profitability of the contrarian trading strategy with $m = 60$ seconds during the 2010 Flash Crash applied to stocks within the S&P 1500 universe. Panel B shows the average daily returns of the contrarian trading strategy with $m = 600$ seconds applied to the same same universe of stocks during the 2007 Quant Meltdown. What is remarkable is the similarity in returns between these two market dislocations, despite the order of magnitude difference in time horizon ($m = 60$ seconds vs. $m = 600$ seconds) and the time-scale over which they last (minutes vs. days).

The SEC and US Commodity Futures Trading Commission (CFTC) joint report (2010) on the Flash Crash states that, even before the crash, May 6th was an unusually turbulent day for markets and that, by early afternoon, negative market sentiment was already affecting the volatility of some securities. The profitability of the short-term mean reversion strategy seen in Panel A of Fig. 4-3 started to rise around 1:45 pm ET in response to the increasing volatility and trading activity. To interpret this pattern, recall that the mean reversion strategy provides immediacy by buying losers and selling winners every m seconds. As positions were being deleveraged, the price for immediacy presumably increased, resulting in higher profits for market making strategies such as ours (recall from Eq. (2.11) that expected returns rise when the speed of mean reversion, parameterized by θ_i, increases). Easley et al. (2011) show that one hour prior to the Flash Crash, the stock market recorded its highest reading of order flow toxicity – a measurement of adverse selection in the marketplace – in recent history. This increase in toxicity may be one factor that induced some

46
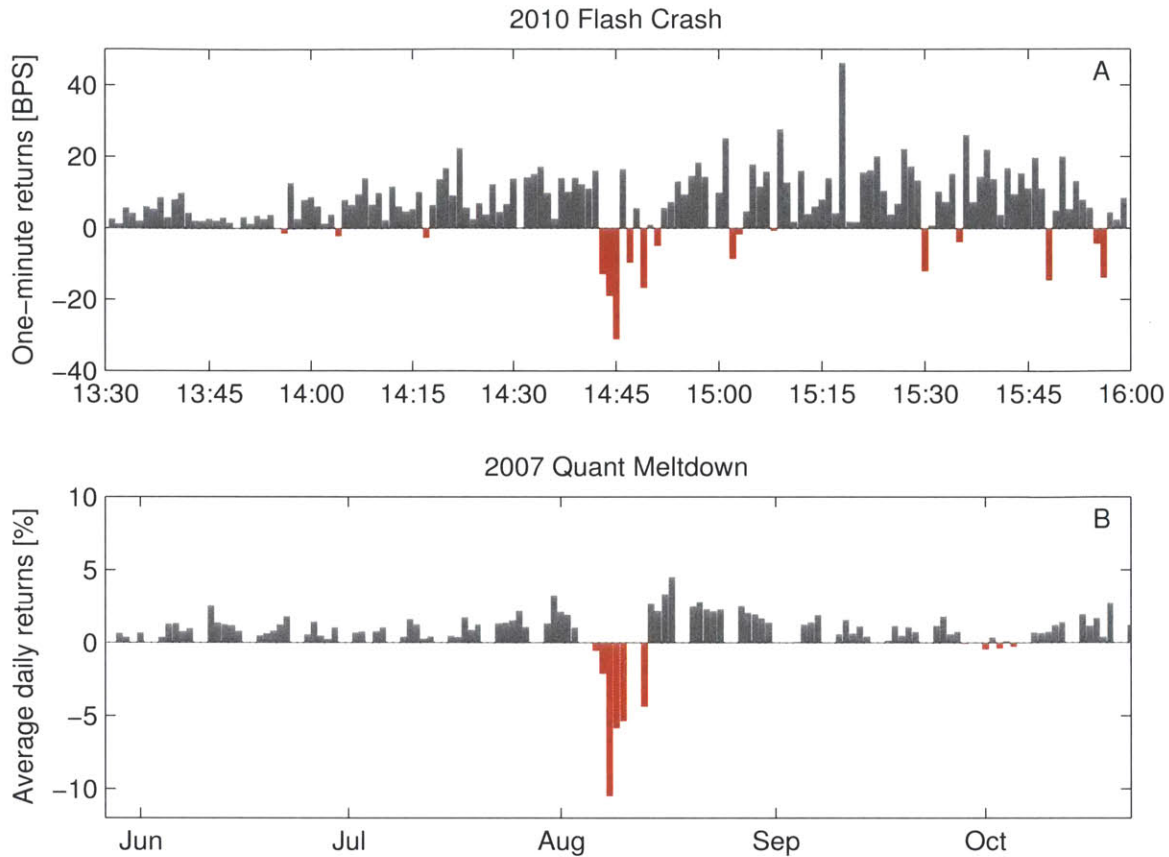
Figure 4-3: Minute-by-minute returns of Lo and MacKinlay's (1990) contrarian trading strategy with $m = 60$ seconds applied to S&P 1500 stocks during the Flash Crash of May 6, 2010 (Panel A). Average daily returns of the same strategy with $m = 600$ seconds applied to S&P 1500 stocks during the Quant Meltdown of August 2007 (Panel B).

participants to withdraw from the market, reducing liquidity. A similar pattern is seen in Panel B in the weeks leading up to August 6, 2007. Khandani and Lo (2011) provide evidence that in the weeks leading up to the Quant Meltdown long/short equity positions were being unwound which would have similarly increased the price for immediacy, inducing higher profits for market making strategies such as ours.

The SEC and CFTC go on to report that at 2:32 pm ET, amid thinning liquidity and unusually high volatility, "a large fundamental trader (a mutual fund complex) initiated a sell program to sell a total of 75,000 E-Mini S&P 500 contracts (valued at approximately $4.1 billion) as a hedge to an existing equity position." Within minutes,

HFT firms and other traders, began quickly buying and re-selling the long futures positions, with an overall effect of driving the E-mini S&P 500 down approximately 3% in just four minutes from 2:41 pm to 2:44 pm which soon spilled over into the equities market. As seen in Panel A of Fig. 4-3, the return of our mean reversion strategy turned sharply negative, which is consistent with the large and sustained negative pressure that was placed on security prices. Again, the same pattern can be found during the Quant Meltdown, however over a longer time horizon and sustained over days rather than minutes. The cause of the sharp decline in this case was likely caused by the unwind of large quant factor subsets causing other hedge funds employing similar strategies to deleverage to mitigate losses, yet inducing the unfortunate effect of prolonging the unwind. While the exact cause of the Flash Crash is still debated, what is clear is that, in both these cases, complex market connections fostered the right conditions to create a positive feedback loop of sustained negative returns. Some form of toxicity reduced liquidity in the marketplace and, while markets were in this fragile state, the right stimulus caused the system to fail.

The joint report states that, at 2:45:33 pm, prices stabilized shortly after trading on the E-Mini was paused for five seconds when the Chicago Mercantile Exchange Stop Logic Functionality was triggered in order to prevent further price declines. During that time buy and sell interest was allowed to recalibrate and a normal price discovery process was restored. In the aftermath of the Flash Crash, contrarian trading strategy profits remained elevated relative to their recent historical average and it was not for a few days that new capital began to flow into the market to take advantage of opportunities created by the dislocation. Similarly, by Tuesday, August 14, 2007 the supply/demand imbalances during the Quant Meltdown returned to relatively normal levels, and market making capital began to flow back to take advantage of the higher profits.

Fig. 4-4 shows the minute-by-minute returns of the contrarian strategy applied to the S&P 1500 stocks for $m = 15$, 30 and 60 seconds. The one-minute returns of the $m < 60$ second trading strategies are calculated by summing their returns within each one-minute interval. Between 2:41 and 2:45:27 pm ET, the S&P 500 SPR exchange
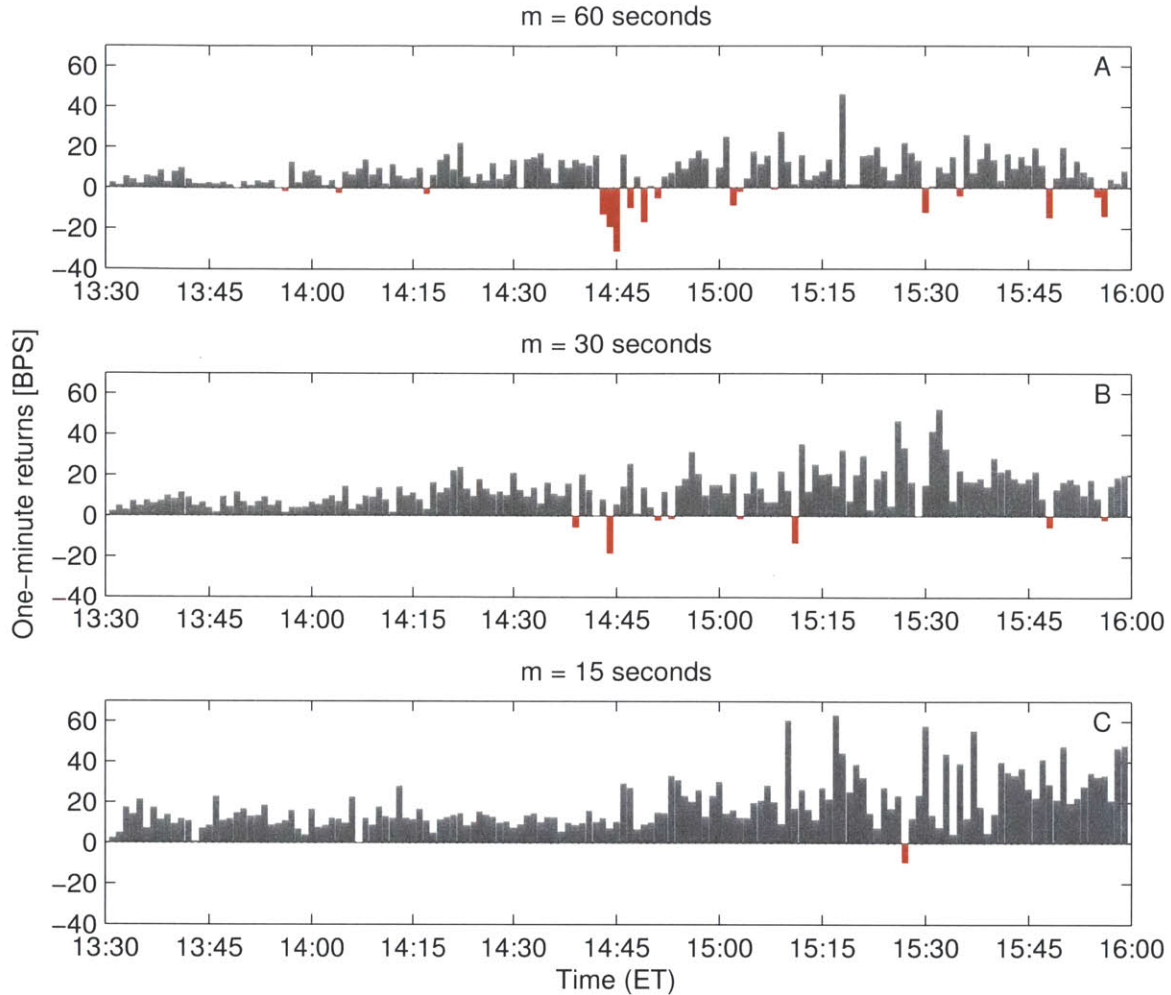
Figure 4-4: Minute-by-minute returns of Lo and MacKinlay's (1990) contrarian trading strategy with $m = 15$, 30 and 60 seconds (Panels A, B and C, respectively) applied to S&P 1500 stocks during the Flash Crash of May 6, 2010.

traded fund suffered a decline of over 6%, the momentum of which was translated into substantial losses for the $m = 60$ second contrarian strategy. The higher frequency contrarian strategies suffered less severely, and in fact, Panel C shows the $m = 15$ second strategy was completely unaffected by the contemporaneous market losses during this period. This observation is consistent with the findings of Kirilenko et al. (2011) who show that, during the 2010 Flash Crash, HFTs, despite their large share of trading volume, appeared to operate at such a high speed that they were not affected by the price level or volatility. Intermediaries, on the other hand, traded on a slower timescale and were caught on the wrong side of the market as prices fell.

Chapter 5

Spectral Analysis

Fig. 4-4 from the previous chapter demonstrates that the contrarian strategy has different risk-reward characteristics across different holding periods. To further understand this observation we develop a spectral analysis technique that can be used to decompose any portfolio's average return into its frequency components. The discussion closely parallels Engle (1974, 1978), Hasbrouck and Sofianos (1993) and Baron, Brogaard and Kirilenko (2012), however we make a novel modification to their analysis to make it applicable to weights and returns. We go on to relate this spectral analysis as an extension of Lo's (2008) active/passive (AP) decomposition.

Fourier transform representations of portfolio returns provide us with a more direct characterization of time horizon effects. Simply put, once two time series are expressed in terms of their frequency components, it becomes possible to decompose their cross-moment into the contributions from different frequencies. Since expected returns can be approximated by the cross-moment of weights and returns, spectral analysis can be applied directly[1]. The Discrete Fourier Transform (DFT) is used to extract the frequency components of the weight and return time series. If the weights and returns are in phase at a given frequency, the contribution that frequency makes to the average portfolio return is positive. If they are out of phase, then that particular frequency's contribution will be negative.

---

[1]This expression is an approximation because it ignores the effect of compounding. The error is typically small, especially for short holding periods

For a stock i held in a portfolio from times $n = 0, ..., N - 1$, the average weighted one-period return can be calculated as:

$$\overline{w_i r_i} = \frac{1}{N} \sum_{n=0}^{N-1} w_i[n] r_i[n], \tag{5.1}$$

where $w_i[n]$ is the weight and $r_i[n]$ is the return of the ith stock at time n. This calculation is equivalent to the one formed from the N-point DFT:

$$\overline{w_i r_i} = \frac{1}{N^2} \sum_{k=0}^{N-1} W_i[k] R_i^*[k] = \frac{1}{N^2} \sum_{k=0}^{N-1} Re(W_i[k] R_i^*[k]) \tag{5.2}$$

where the asterisk (*) denotes complex conjugation, Re denotes taking the real component of a complex number, and $W_i[k]$ and $R_i[k]$ are the N-point DFT coefficients of $w_i[n]$ and $r_i[n]$ respectively. The last equality follows because the imaginary components sum to 0 (see Appendix B.1 for details). If we consider M stocks, the average one-period portfolio return is simply:

$$\overline{r_p} = \sum_{i=1}^{M} \overline{w_i r_i} = \frac{1}{N^2} \sum_{i=1}^{M} \sum_{k=0}^{N-1} Re(W_i[k] R_i^*[k]). \tag{5.3}$$

In this form, the contribution to the average one-period portfolio return by the frequency $f_k = k/(NT_s)$, where $k \in 0, ..., N - 1$ and T_s is the length of time of one period, is clearly visible. Note that $Re(W_i[k] R_i^*[k])$, for $k \in 1, ..., N - 1$, is symmetric about $k = N/2$ and that values of k closer to $N/2$ correspond to higher frequencies. The value at $k = 0$ is the lowest (DC) frequency and, as we show in Appendix B.3, is computationally equivalent to the "Passive Component" defined in Lo (2008). Moreover, the higher frequency components correspond to a further decomposition of Lo's "active component", allowing us to visualize the profitability of different time horizon strategies employed by a portfolio manager. Similar to the "active/passive" (AP) decomposition, if the weights or returns are being sampled at a lower frequency than the weights are being changed, then the decomposition will be biased due to a phenomenon known as aliasing.

5.1   Some Numerical Examples

In this section, we apply spectral analysis to the numerical example portfolios described in Lo (2008) to provide intuition for the frequency decomposition and to illustrate its connection with Lo's active and passive components. First, consider a portfolio of two assets, one which yields a monthly return that alternates between 1% and 2% (Asset 1) and the other which yields a fixed monthly return of 0.15% (Asset 2). Let the weights of this portfolio, called A1, be given by 75% in Asset 1 and 25% in Asset 2. Table 5.1 illustrates the dynamics of this portfolio over a 12-month period, where the expected return of the portfolio is 1.1625% per month, all of which is due to the DC frequency ($k = 0$).

Table 5.1: The expected return of a constant portfolio depends only on the DC ($k = 0$) component.

Month	w_1	r_1	w_2	r_2	r_p
			Strategy A1		
1	75%	1.00%	25%	0.15%	0.7875%
2	75%	2.00%	25%	0.15%	1.5375%
3	75%	1.00%	25%	0.15%	0.7875%
4	75%	2.00%	25%	0.15%	1.5375%
5	75%	1.00%	25%	0.15%	0.7875%
6	75%	2.00%	25%	0.15%	1.5375%
7	75%	1.00%	25%	0.15%	0.7875%
8	75%	2.00%	25%	0.15%	1.5375%
9	75%	1.00%	25%	0.15%	0.7875%
10	75%	2.00%	25%	0.15%	1.5375%
11	75%	1.00%	25%	0.15%	0.7875%
12	75%	2.00%	25%	0.15%	1.5375%
Mean:	75%	1.50%	25%	0.15%	**1.1625%**

		Frequency decomposition of $\overline{r_p}$				
$k = 0$	$k = 1, 11$	$k = 2, 10$	$k = 3, 9$	$k = 4, 8$	$k = 5, 7$	$k = 6$
1.1625%	0%	0%	0%	0%	0%	0%

Now consider portfolio A2, which differs from A1 only in that the portfolio weight for Asset 1 alternates between 50% and 100% in phase with Asset 1's returns which alternates between 1% and 2% (see Table 5.2). In this case, the total expected return is 1.2875% per month, of which 0.1250% is due to the positive correlation between the

53

portfolio weight for Asset 1 and its return at the shortest-time horizon (i.e., highest frequency).

Table 5.2: The dynamics of the portfolio weights are positively correlated with returns at the shortest time horizion, which adds value to the portfolio and yields a positive contribution from the highest frequency ($k = 6$).

Month	w_1	r_1	w_2	r_2	r_p
			Strategy A2		
1	50%	1.00%	50%	0.15%	0.5750%
2	100%	2.00%	0%	0.15%	2.0000%
3	50%	1.00%	50%	0.15%	0.5750%
4	100%	2.00%	0%	0.15%	2.0000%
5	50%	1.00%	50%	0.15%	0.5750%
6	100%	2.00%	0%	0.15%	2.0000%
7	50%	1.00%	50%	0.15%	0.5750%
8	100%	2.00%	0%	0.15%	2.0000%
9	50%	1.00%	50%	0.15%	0.5750%
10	100%	2.00%	0%	0.15%	2.0000%
11	50%	1.00%	50%	0.15%	0.5750%
12	100%	2.00%	0%	0.15%	2.0000%
Mean:	75%	1.50%	25%	0.15%	**1.2875%**

Frequency decomposition of $\overline{r_p}$

$k = 0$	$k = 1, 11$	$k = 2, 10$	$k = 3, 9$	$k = 4, 8$	$k = 5, 7$	$k = 6$
1.1625%	0%	0%	0%	0%	0%	0.1250%

Finally, consider a third portfolio A3 which also has alternating weights for Asset 1, but exactly out of phase with Asset 1's returns. When the return is 1%, the portfolio weight is 100%, and when the return is 2%, the portfolio weight is 50%. Table 5.3 confirms that this is counterproductive as Portfolio A3 loses 0.1250% per month from its highest frequency component, and its total expected return is only 1.0375%.

Note that in all three cases, the DC components ($k = 0$) are identical at 1.1625% per month because the average weight for each asset is the same across all three port-folios. The only differences among A1, A2, and A3 are the dynamics of the portfolio weights at the shortest time horizon, and these differences give rise to different values for the highest frequency component. As shown in Appendix B.3, the contribution from the DC frequency is computationally equivalent to Lo's passive component and

54

Table 5.3: The dynamics of the portfolio weights are negatively correlated with returns at the shortest time horizion, which subtracts value from the portfolio and yields a negative contribution from the highest frequency ($k = 6$).

Month	w_1	r_1	w_2	r_2	r_p
		Strategy A3			
1	100%	1.00%	0%	0.15%	1.0000%
2	50%	2.00%	50%	0.15%	1.0750%
3	100%	1.00%	0%	0.15%	1.0000%
4	50%	2.00%	50%	0.15%	1.0750%
5	100%	1.00%	0%	0.15%	1.0000%
6	50%	2.00%	50%	0.15%	1.0750%
7	100%	1.00%	0%	0.15%	1.0000%
8	50%	2.00%	50%	0.15%	1.0750%
9	100%	1.00%	0%	0.15%	1.0000%
10	50%	2.00%	50%	0.15%	1.0750%
11	100%	1.00%	0%	0.15%	1.0000%
12	50%	2.00%	50%	0.15%	1.0750%
Mean:	75%	1.50%	25%	0.15%	**1.0375%**

Frequency decomposition of $\overline{r_p}$

$k = 0$	$k = 1, 11$	$k = 2, 10$	$k = 3, 9$	$k = 4, 8$	$k = 5, 7$	$k = 6$
1.1625%	0%	0%	0%	0%	0%	-0.1250%

contributions from higher frequencies ($k > 0$) sum to Lo's active component. These higher frequency contributions can be interpreted as the portion of the active component that arises from a given time horizon.

5.2 Analytical Result of a Simple MA(1) Model

Now suppose that stock returns are i.i.d and satisfy the following simple MA(1) model:

$$r_i[n] = \varepsilon_i[n] + \lambda\varepsilon_i[n-1] \tag{5.4}$$

where the $\varepsilon_i[n]$ are serially and cross-sectionally uncorrelated white-noise random variables with variance σ^2.

So far we have shown how to calculate the spectral representation of a portfolio's average return for a finite sample of observations. Here we calculate the population

analog, called the cross-spectrum, for the simple stochastic process described in Eq. (5.4). If we simplify the contrarian trading strategy described in Eq. (2.3) so that we do not normalize by $I[n]^2$ and set $k = 1$, then the expected one-period portfolio return is:

$$E[r_p] = -\lambda\sigma^2(1 - \frac{1}{M}) \qquad (5.5)$$

We see that profits are proportional to both the mean reversion factor λ as well as the variance of the underlying stock returns. To construct the cross-spectrum for each security, $\Psi_{w_i r_i}(\omega)$, we take the Fourier transform of the cross-covariance function:

$$\Psi_{w_i r_i}(\omega) = \sum_{l=-\infty}^{\infty} Cov(r_i[n], w_i[n-l])e^{-j\omega l} = -\frac{\sigma^2}{M}(1 - \frac{1}{M})(\lambda e^{j2\omega} + (1 + \lambda^2)e^{jw} + \lambda).$$

$$(5.6)$$

As with the sample cross-spectrum, we notice that the imaginary components at complementary frequencies (e.g., $\omega = \{-\pi/4, \pi/4\}$) are equal in magnitude and opposite in sign, so we can define the spectral representation of each security's expected one-period weighted return as the real component of Eq. (5.6):

$$Re(\Psi_{w_i r_i}(\omega)) = -\frac{\sigma^2}{M}(1 - \frac{1}{M})(\lambda cos(2\omega) + (1 + \lambda^2)cos(\omega) + \lambda). \qquad (5.7)$$

The expected one-period portfolio return can be recovered by summing over the expected weighted return of each security which is itself equivalent to the average value of Eq. (5.7) over a continuous interval of 2π.

$$E[r_p] = \sum_{i=1}^{M} \frac{1}{2\pi} \int_{-\pi}^{\pi} Re(\Psi_{w_i r_i}(\omega)) \, d\omega = -\lambda\sigma^2(1 - \frac{1}{M}), \qquad (5.8)$$

Note that the general form of Eq. (5.8) is the population analog to Eq. (5.3).

Fig. 5-1 plots the real component of the cross-spectrum for $\lambda = -0.4\sigma$ (mean reversion), 0 (serially uncorrelated) and 0.4σ (momentum). In Panel A we can understand why the lower frequency contributions are negative by recalling that the DC

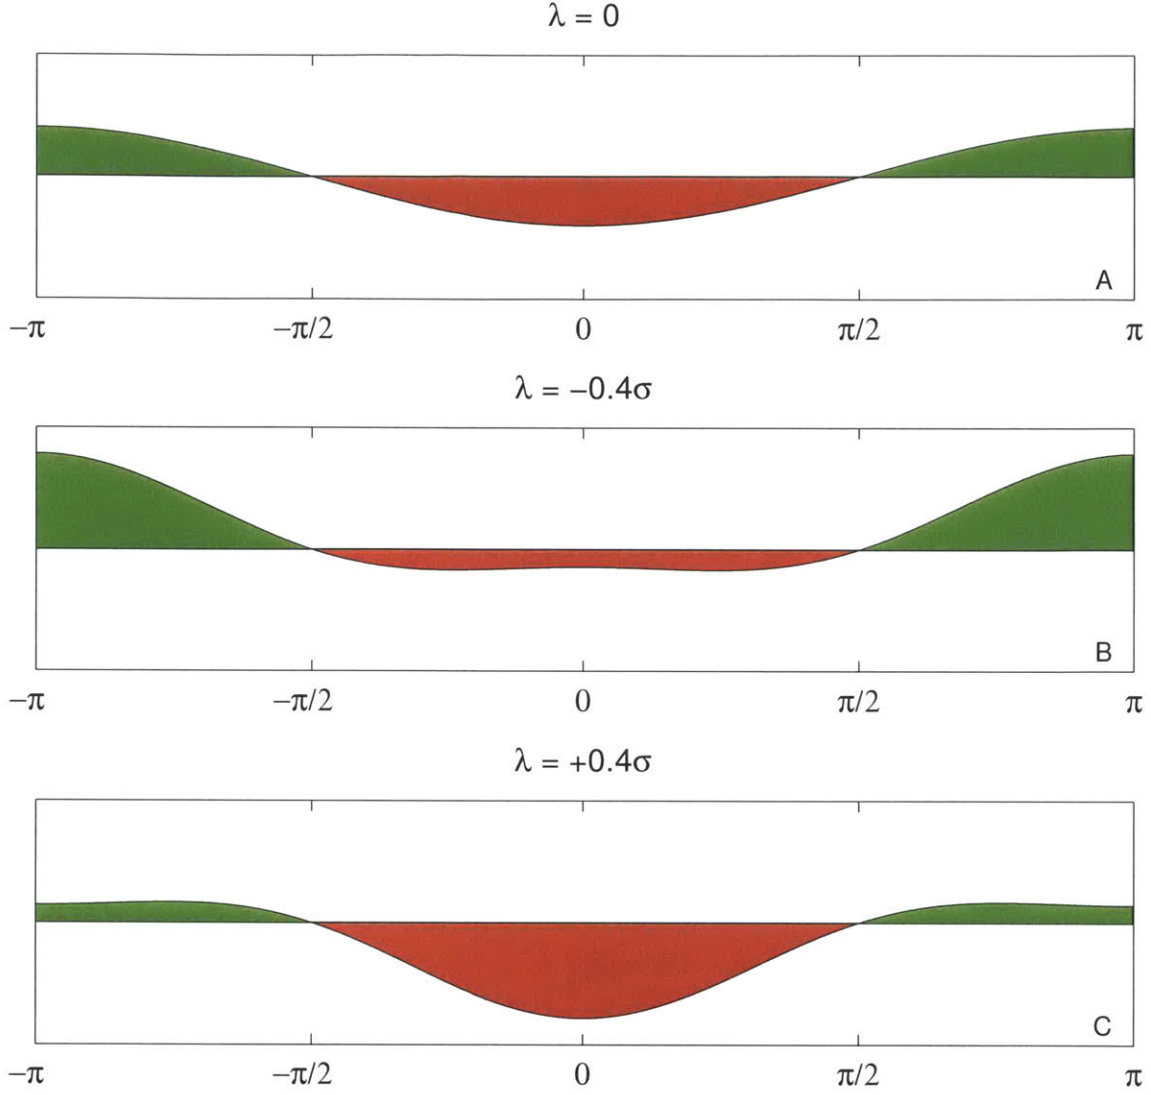[2]This is done to make the analytical result tractable and exact.

Figure 5-1: Spectral analysis of the contrarian trading strategy applied to serially un-correlated (Panel A), mean reversion (Panel B) and momentum (Panel C) realizations of Eq. (5.4).

component of Eq. (5.3) is equal to $\sum_{i=1}^{M} \overline{w_i} \, \overline{r_i}$. Since returns are uncorrelated, if $\overline{r_i}$ is negative (positive), then $\overline{w_i}$ will tend to be positive (negative) due to the contrarian aspect of the trading strategy. Their product will therefore be negative. Similarly, the highest frequency's contribution is positive because, if $r_i[n]$ has significant power in that frequency (i.e., it randomly has reversions in that frequency) then, because of the nature of the contrarian trading strategy, so will $w_i[n]$. These high-frequencies components [see Eq. (B.26)] will tend to be in phase leading to an overall positive

contribution. We also notice that the positive contributions from the high frequencies and the negative contributions from the low frequencies cancel and so, for this model of serially and cross-sectionally uncorrelated returns, the average profitability of the contrarian trading strategy is 0.

Panels B and C of Fig. 5-1 show spectral analysis applied to contrarian trading profits for mean reverting and momentum returns. For the mean reversion case we notice that, relative to the serially uncorrelated case, both the lowest and highest frequencies are more profitable, while the middle frequencies remain approximately unchanged. Furthermore, the integral over all frequency components, and therefore the overall profitability, is positive. The momentum case is the opposite. Relative to the serially uncorrelated case, both the lowest and highest frequencies are less profitable, the middle frequencies remain approximately unchanged, and the integral over all frequency components is negative.

Now suppose that in each of the above cases, the standard deviation of the white noise variable doubles (e.g., a widening of the inside spread). In this case, the contribution to the average portfolio return from each frequency quadruples.

5.3 Diversification Across Time Horizons

For quantitative strategies like the contrarian trading strategy where we are investing at a particular time horizon, we may consider diversifying not only across assets but also across the frequency components of trading strategies. We know from chapter 4 that, in general, operating at a higher frequency earns higher returns and therefore time horizon diversification will result in smaller returns. What we learned in section 4.1 was that market dislocations can be isolated to certain frequency bands, and therefore, contrarian trading profits across multiple frequencies need not be strongly correlated during market dislocations. Since these time horizon specific strategies can be run contemporaneously, they can be viewed as separate assets with varying risk-reward characteristics and correlations. One could then optimize the expected profit of a portfolio that held these frequency dependent assets for a given variance

and risk-related measures.

Fig. 5-2 shows the frequency decomposition of each strategy's profits averaged over the entire two decade time period. We notice that most of the profits for each strategy are isolated to a band around the targeted frequency, and therefore are most sensitive to changes in the idiosyncratic mean reversion at that frequency. Therefore, as was demonstrated for the Flash Crash in section 4.1, an increase in momentum at one frequency could substantially impact the profitability of one trading strategy while leaving another unaffected.
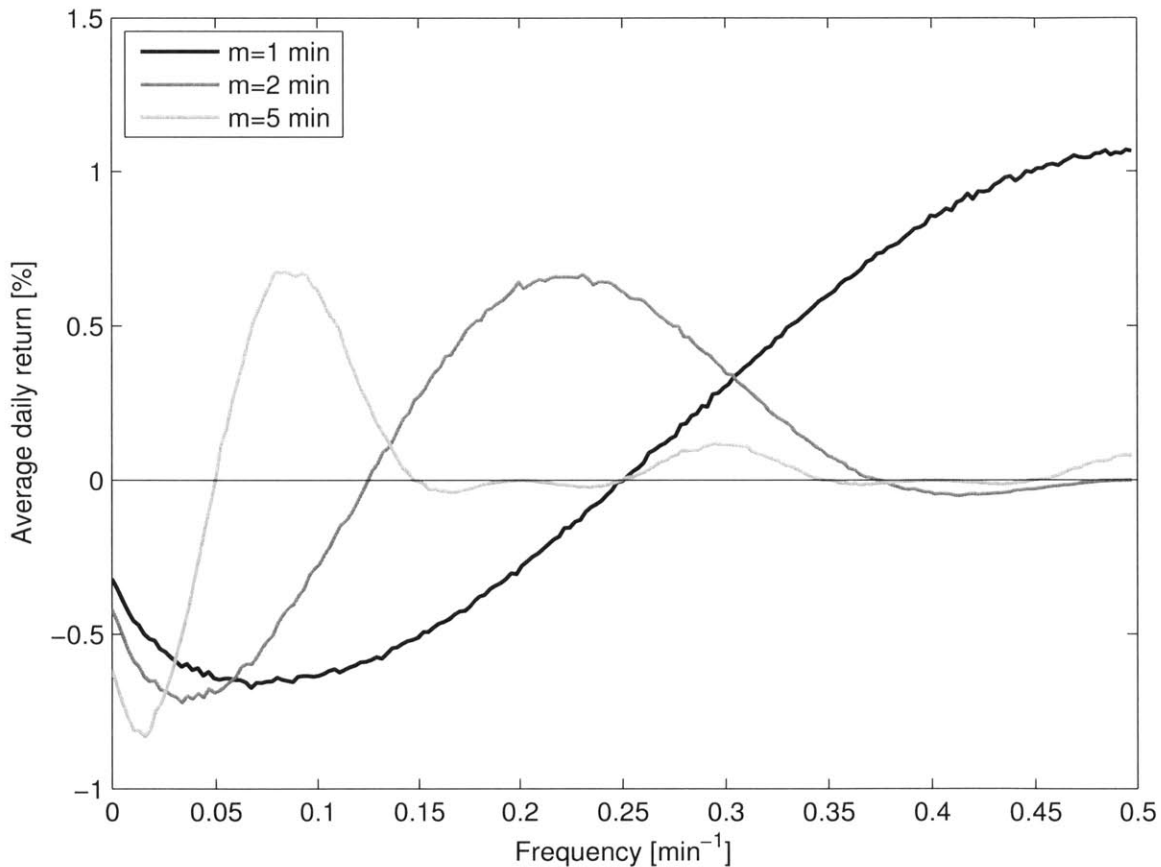

Figure 5-2: Frequency contributions to yearly average of daily profits for various contrarian strategies with different time horizons averaged over 1994 to 2012.

If one considers each frequency's (or frequency band's) contribution an asset with its own time-varying mean, variance and covariance with other frequencies within the same and across other trading strategies, then each strategy could be thought of as representing a portfolio of "frequency assets". These parameters could then be

estimated using the spectral analysis technique described herein, however to increase efficiency and accuracy a method of directly estimating the variances and covariances of these frequency assets would have to be developed. One could then apply the optimization techniques from modern portfolio theory to combine various trading strategies to efficiently manage risk at various time horizons. Another method would be to consider each trading strategy as a separate time-series of returns and use spectral analysis to estimate how the profitability and variance at a given frequency band would change as a function of portfolio allocation. These diversification across frequency applications require further analysis and will be the focus of future research.

Chapter 6

Conclusion

In this thesis we applied time and frequency domain analyses to contrarian trading profits to help untangle the role regulations, market environment and technology played in shaping recent market events. In particular, this strategy enabled us to study the profits of liquidity provision since, by buying recent losers and selling recent winners, we were effectively able to capture profits from the bid-ask spread by transforming the negative autocorrelation from market making activity (e.g., the bid-ask bounce) into positive returns. Moreover, by studying this profitability over multiple decades and time horizons we were able to characterize gradual changes to the market microstructure from increased competition to traditional market makers and also rapid shifts in price dynamics from market dislocations and new regulations.

We began by studying Lo and MacKinlay's (1990) contrarian trading strategy when applied to the daily returns of exchange sorted portfolio subsets. The large rise and subsequent decline in profitability when the strategy was applied to the subset of NASDAQ NMS securities throughout the 1980s and 1990s was in sharp contrast to the flat, near-zero returns delineated by the NYSE subset. The source of these profits stemmed from the percentage bid-ask spread which widened and narrowed in tandem with the NMS profits. We found evidence that market makers were implicitly colluding, across all sizes of NMS stocks, to offer only even-eighth bid-ask quotes thereby maintaing at least a $0.25 spread in an effort to increase profits at the expense of larger transaction costs for investors. Competitive pressure from

61

the entry of non-traditional liquidity providers in the form of algorithmic traders willing to offer odd-eighth quotes to earn a greater share of order flow seems to have reduced the profitability of these inefficient practices. New regulations requiring stocks to be traded in sixteenths and then hundredths of a dollar also helped to narrow inside spreads – however, the rate of adoption of new price levels was slower for the NMS compared to the NYSE consistent with the premise of greater market maker inefficiency.

Next, to improve temporal resolution and study changes to markets as traders began to operate at faster time-scales, we applied the contrarian strategy to minute-level intraday holding periods again sorting securities by their primary exchange. For both NASDAQ and NYSE subsets there was a general downward trend in the profitability of our mean reversion strategies driven by increasing levels of liquidity. Increased market making competition, advances in algorithmic trading and the introduction of smaller minimum price variations all contributed to the gradual decline in profitability which was frequently interrupted by liquidity shocks and market dislocations. By reducing the length of the holding period we found that profitability increased and sensitivity to market instability decreased providing strong incentives for traders to move into higher frequency domains.

Increasing the focal strength of our time domain analysis by another order of magnitude, we analyzed the minute-by-minute contrarian trading profits during the 2010 Flash Crash and compared it to the 10 min holding period returns of the same strategy during the 2007 Quant Meltdown. The similarity of these time-series led us to conclude that while the event-specific details and time-scales of these two dislocations were different, the nature of their underlying causes were similar. In both cases, complex market connections fostered the right conditions to create a positive feedback loop of sustained negative returns. Some form of toxicity reduced liquidity in the marketplace and, while markets were in this fragile state, the right stimulus caused the system to fail. As prices stabilized, liquidity providers that were able to remain solvent earned large profits and eventually market making capital flowed back into the system. By decreasing the holding period of the contrarian strategy from minutes

to seconds we found that profitability was no longer adversely affected by the sudden and widespread market losses that occurred during the Flash Crash. This suggests that very high frequency traders were unaffected by the market downturn while slower traders were been caught on the wrong side of the market as prices fell.

This last observation led us to consider a frequency domain analysis of profits. We developed a spectral analysis technique that extended Lo's (2008) active/passive decomposition into frequency specific components and enabled us to characterize the profitability of a trading strategy at different time horizons. An analytical analysis of a simple MA(1), i.i.d returns model showed that the profitability of the contrarian strategy was greatest around its holding period and weakest around the DC (i.e. long-term, passive) frequency. This suggests that turbulence or momentum in one frequency band may substantially affect the contrarian strategies whose holding periods target those frequencies, while leaving shorter holding period strategies relatively unaffected. This presents a plethora of diversification across frequencies applications that we hope to explore in future work.

Appendix A

Sampling Theory for $\hat{\tilde{C}}_k$, $\hat{\tilde{O}}_k$, and $\hat{\sigma}^2(\tilde{\mu})$

To derive the sampling theory for the estimators $\hat{\tilde{C}}_k$ and $\hat{\tilde{O}}_k$, we reexpress them as averages of artificial time series and then apply standard asymptotic theory to those averages. For details on the required assumptions we refer the reader to Appendix 2 to in Lo and MacKinlay (1990).

Let $\tilde{C}_k[n]$ and $\tilde{O}_k[n]$ be the following two time series:

$$\tilde{C}_k[n] = r_m[n]\tilde{r}_m[n-k] - \hat{\mu}_m\hat{\tilde{\mu}}_m - \frac{1}{M^2}\sum_{i=1}^{M}(r_i[n]\tilde{r}_i[n-k] - \hat{\mu}_i\hat{\tilde{\mu}}_i) \tag{A.1}$$

$$\tilde{O}_k[n] = -(\frac{M-1}{M^2})\sum_{i=1}^{M}(r_i[n]\tilde{r}_i[n-k] - \hat{\mu}_i\hat{\tilde{\mu}}_i) \tag{A.2}$$

where $\hat{\mu}_i$ and $\hat{\mu}_m$ are the usual sample means of the returns to security i and the equal-weighted market index, respectively, and $\hat{\tilde{\mu}}_i$ and $\hat{\tilde{\mu}}_m$ are their normalized return $(\tilde{r}_i[n] = r_i[n]/I[n])$ counterparts. Then the estimators $\hat{\tilde{C}}_k$, $\hat{\tilde{O}}_k$, and $\hat{\sigma}^2(\tilde{\mu})$ are given

by,

$$\hat{\tilde{C}}_k = \frac{1}{N-(k+1)} \sum_{n=k+1}^{N-1} \tilde{C}_k[n] \qquad (A.3)$$

$$\hat{\tilde{O}}_k = \frac{1}{N-(k+1)} \sum_{n=k+1}^{N-1} \tilde{O}_k[n] \qquad (A.4)$$

$$\hat{\sigma}^2(\tilde{\mu}) = \frac{1}{M} \sum_{i=1}^{M} (\hat{\mu}_i - \hat{\mu}_m)(\hat{\tilde{\mu}}_i - \hat{\tilde{\mu}}_m). \qquad (A.5)$$

Appendix B

Spectral Analysis Relations

In this Appendix, we provide proofs for the main spectral analysis relations in the thesis.

B.1 Frequency Decomposition of the Cross-Moment

If $x_1[n]$ and $x_2[n]$ are two real-valued time-series defined on $n = 0, ..., N - 1$, show that,

$$\frac{1}{N} \sum_{n=0}^{N-1} x_1[n]x_2[n] = \frac{1}{N^2} \sum_{k=0}^{N-1} X_1[k]X_2^*[k] = \frac{1}{N^2} \sum_{k=0}^{N-1} Re(X_1[k]X_2^*[k]) \qquad \text{(B.1)}$$

where $X_1[k]$ and $X_2[k]$ are the N-point DFT coefficients of $x_1[n]$ and $x_2[n]$ respectively. Specifically,

$$X_i[k] = \sum_{n=0}^{N-1} x_i[n]e^{-j\frac{2\pi kn}{N}}, \qquad 0 \le k \le N - 1 \qquad \text{(B.2)}$$

$$x_i[n] = \frac{1}{N} \sum_{k=0}^{N-1} X_i[k]e^{j\frac{2\pi kn}{N}}, \qquad 0 \le n \le N - 1. \qquad \text{(B.3)}$$

Proof: Let $x_3[n] = x_2[((-n))_N]$, for $n = 0, ..., N - 1$, where the notation $((k))_N$ represents the modulus of k after division by N. Also let $y[n]$ be the N-point circular

convolution of $x_1[n]$ and $x_3[n]$, or more specifically,

$$y[n] = \sum_{m=0}^{N-1} x_1[m]x_3[((n-m))_N].\tag{B.4}$$

Evaluating $y[n]$ at $n = 0$ we find that,

$$y[0] = \sum_{m=0}^{N-1} x_1[m]x_3[((0-m))_N]\tag{B.5}$$

$$= \sum_{m=0}^{N-1} x_1[m]x_2[m].\tag{B.6}$$

Similarly, in the frequency domain, we find that,

$$Y[k] = X_1[k]X_3[k]\tag{B.7}$$

$$= X_1[k]X_2^*[k],\tag{B.8}$$

where in the last step we used the fact that $x_2[n]$ is real and the N-point DFT of $x_i^*[((-n))_N]$ is equal to $X_i^*[k]$. Using Eq. (B.3), we see that,

$$y[0] = \frac{1}{N}\sum_{k=0}^{N-1} Y[k]\tag{B.9}$$

$$= \frac{1}{N}\sum_{k=0}^{N-1} X_1[k]X_2^*[k]\tag{B.10}$$

$$= \frac{1}{N}\sum_{k=0}^{N-1} Re(X_1[k]X_2^*[k]).\tag{B.11}$$

Where in the last step we noticed that $y[0]$ is real and so the imaginary components of $X_1[k]X_2^*[k]$ must sum to 0. Finally, we achieve the desired result by equating Eq. (B.6) and Eq. (B.11), and multiplying both sides by an extra factor of N^{-1}.

B.2   Symmetry in Spectral Analysis

Show that $Re(X_1[k]X_2^*[k])$, for $k = 1, ..., N-1$, is symmetric about $N/2$

Proof: From Eq. (B.2) we know that for $k = 0, ..., N - 1$,

$$X_1[N - k] = \sum_{n=0}^{N-1} x_1[n]e^{-j\frac{2\pi(N-k)n}{N}} \tag{B.12}$$

$$X_1^*[N - k] = \sum_{n=0}^{N-1} x_1[n]e^{j\frac{2\pi(N-k)n}{N}} \tag{B.13}$$

$$= \sum_{n=0}^{N-1} x_1[n]e^{-j\frac{2\pi kn}{N}} \tag{B.14}$$

$$= X_1[k], \tag{B.15}$$

where, is the second step, we used the fact that the complex conjugate is distributive under complex addition. Similarly, we find that $X_2^*[k] = X_2[N - k]$ for $k = 0, ..., N - 1$. Putting it all together and using the fact that the complex conjugate is also distributive over complex multiplication we find that for $0 \leq k \leq N - 1$,

$$X_1[k]X_2^*[k] = (X_1[N - k]X_2^*[N - k])^* \tag{B.16}$$

$$Re(X_1[k]X_2^*[k]) = Re(X_1[N - k]X_2^*[N - k]). \tag{B.17}$$

From Eq. (B.17) it becomes clear that $Re(X_1[k]X_2^*[k])$, for $k = 1, ..., N - 1$, is symmetric about $N/2$. Therefore it is sufficient to consider only frequencies where $0 \leq k \leq \lfloor \frac{N}{2} \rfloor$ when performing spectral analysis.

B.3 Time Horizon Extension of the AP Decomposition

Show that Lo's (2008) passive component (ν_p) is computationally equivalent to the DC component from the spectral analysis and that his active component (δ_p) is equivalent to the sum of contributions from higher frequencies.

Lo (2008) shows how to decompose the average portfolio return ($\overline{r_p}$) into an active

69

component (δ_p) and a passive component (ν_p). The value of δ_p is attributed to active management, as it estimates the profitability of the portfolio manager's conscious decision to buy, sell, or avoid a security by aggregating the sample covariances between the portfolio weights, $w_i[n]$, and security returns, $r_i[n]$. In particular, if a manager has, on average, positive weights when security returns are positive and negative weights when returns are negative, this implies positive covariances between portfolio weights and returns, and will have a positive impact on the portfolio's average return. In effect, the covariance term captures the manager's timing ability, asset by asset. The non-DC component contributions to trading profits decompose this covariance term further, by calculating the profitability of different time horizons employed by a trading strategy. Fourier analysis first deconstructs the weights and returns into their various frequency components. At each frequency, if the weights and returns are in phase (i.e., peaks and valleys coincide), then that time horizon's contribution to the average portfolio return will be positive. If the two signals are out of phase, then that particular frequency's contribution will be negative. As we show show in Eq. (B.24), the sum of all contributions made by non-DC frequencies across all assets is δ_p.

The other source of profitability that must be added to δ_p to get the average portfolio return is ν_p, which has nothing to do with asset timing, but arises instead from the manager's average position in a security. Lo labels this contribution "passive" and, as shown in Eq. (B.20), it is equivalent to the DC component's contribution to trading profits. Specifically,

$$\nu_p = \sum_{i=1}^{M} \overline{w_i}\,\overline{r_i}, \tag{B.18}$$

where

$$\overline{w_i} = \frac{1}{N} \sum_{n=0}^{N-1} w_i[n], \qquad \overline{r_i} = \frac{1}{N} \sum_{n=0}^{N-1} r_i[n]. \tag{B.19}$$

Similarly, in the frequency domain, the contribution to the average portfolio return from the DC component can be calculated as,

$$\frac{1}{N^2} \sum_{i=1}^{M} Re(W_i[0]R_i^*[0]) = \sum_{i=1}^{M} \overline{w_i}\,\overline{r_i} = \nu_p, \tag{B.20}$$

where we used Eq. (B.2) to find that,

$$W_i[0] = \sum_{n=0}^{N-1} w_i[n] = N\overline{w_i}, \qquad R_i^*[0] = \sum_{n=0}^{N-1} r_i[n] = N\overline{r_i}. \qquad (B.21)$$

It is now trivial to prove the active component portion of this proposition,

$$\delta_p = \overline{r_p} - \nu_p \qquad (B.22)$$

$$= \frac{1}{N^2} \sum_{k=0}^{N-1} \sum_{i=1}^{M} Re(W_i[k]R_i^*[k]) - \nu_p \qquad (B.23)$$

$$= \frac{1}{N^2} \sum_{k=1}^{N-1} \sum_{i=1}^{M} Re(W_i[k]R_i^*[k]) \qquad (B.24)$$

We end this section by providing some intuition for the time horizon interpretation of the frequency decomposition and its relation to Lo's active component. If we take the inverse transform of $W_i[k]$ and $R_i[k]$, restricted to a single frequency $\tilde{k}$, and denote these single frequency reconstructions as $\tilde{w}_i[n]$ and $\tilde{r}_i[n]$, then the sum of their real and imaginary component cross-moments is equal to that frequency's profitability. Specifically if,

$$\tilde{w}_i[n] = \frac{1}{N} W_i[\tilde{k}]e^{j\frac{2\pi \tilde{k}n}{N}}, \qquad \tilde{r}_i[n] = \frac{1}{N} R_i[\tilde{k}]e^{j\frac{2\pi \tilde{k}n}{N}}, \qquad 0 \le n \le N-1, \qquad (B.25)$$

then,

$$\frac{1}{N} \sum_{n=0}^{N-1} Re(\tilde{w}_i[n])Re(\tilde{r}_i[n]) + Im(\tilde{w}_i[n])Im(\tilde{r}_i[n]) = \frac{1}{N^2} Re(W_i[\tilde{k}]R_i^*[\tilde{k}]). \qquad (B.26)$$

Eq. (B.26) explicitly shows that the contribution from the kth frequency is equivalent to the cross-moment between weights and returns restricted to that time horizon. This interpretation is analogous to that of Lo's AP decomposition and can be considered the restricted frequency version of Eq. (B.1). Finally, we note that this analysis can be restricted to a subset of frequencies to determine the profitability of a band of time horizons. If the kth frequency is included in the band then so should the

$N - k$th frequency since they have the same time horizon (remember that the highest frequency is at N/2). Furthermore, their contributions complement each other since they are complex conjugates and their imaginary components cancel (see Appendix B.2 for more details).

References

Atkins, A. B., & Dyl, E. A. (1990). Price reversals, bid-ask spreads, and market efficiency. *Journal of Financial and Quantitative Analysis*, 25(04), 535-547.

Baron, M., Brogaard, J., & Kirilenko, A. (2012). The trading profits of high frequency traders. *Working Paper*.

Battalio, R. H., Hatch, B., & Jennings, R. (1997). SOES trading and market volatility. *Journal of Financial and Quantitative Analysis*, 32(2).

Bremer, M., & Sweeney, R. J. (1991). The Reversal of Large Stock Price Decreases. *The Journal of Finance*, 46(2), 747-754.

Brown, K. C., Harlow, W. V., & Tinic, S. M. (1988). Risk aversion, uncertain information, and market efficiency. *Journal of Financial Economics*, 22(2), 355-385.

Capgemini. (2012). High Frequency Trading: Evolution and the Future.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of finance*, 52(1), 57-82.

Chordia, T., Roll, R., & Subrahmanyam, A. (2001). Market liquidity and trading activity. *The Journal of Finance*, 56(2), 501-530.

Christie, W. G., & Schultz, P. H. (1994). Why do NASDAQ Market Makers Avoid Odd-Eighth Quotes?. *The Journal of Finance*, 49(5), 1813-1840.

Christie, W. G., Harris, J. H., & Schultz, P. H. (1994). Why Did NASDAQ Market Makers Stop Avoiding Odd-Eighth Quotes?. *The Journal of Finance*,49(5), 1841-1860.

Cox, D. R., & Peterson, D. R. (1994). Stock Returns following Large One?Day Declines: Evidence on Short?Term Reversals and Longer?Term Performance. *The Journal of Finance*, 49(1), 255-267.

De Bondt, W. F., & Thaler, R. (1985). Does the stock market overreact?. *The Journal of finance*, 40(3), 793-805.

Engle, R. F. (1974). Band spectrum regression. *International Economic Review*,15(1), 1-11.

Engle, R. F. (1978). Testing price equations for stability across spectral frequency bands. *Econometrica: Journal of the Econometric Society*, 869-881.

Harris, J. H., & Schultz, P. H. (1997). The importance of firm quotes and rapid executions: Evidence from the January 1994 SOES rules changes. *Journal of Financial Economics*, 45(1), 135-166.

Harris, J. H., & Schultz, P. H. (1998). The trading profits of SOES bandits. *Journal of Financial Economics*, 50(1), 39-62.

Hasbrouck, J., & Sofianos, G. (1993). The trades of market makers: An empirical analysis of NYSE specialists. *The Journal of Finance*, 48(5), 1565-1593.

Jagannathan, R., Malakhov, A., & Novikov, D. (2010). Do hot hands exist among hedge fund managers? An empirical evaluation. *The Journal of Finance*, 65(1), 217-255.

Khandani, A. E., & Lo, A. W. (2007). What happened to the quants in August 2007?. *Journal of Investment Management*, 5(4), 5-54.

Khandani, A. E., & Lo, A. W. (2011). What happened to the quants in August 2007? Evidence from factors and transactions data. *Journal of Financial Markets*, 14(1), 1-46.

Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2011). The flash crash: The impact of high frequency trading on an electronic market. Manuscript, U of Maryland.

Kothare, M., & Laux, P. A. (1995). Trading costs and the trading systems for NASDAQ stocks. *Financial Analysts Journal*, 42-53.

Lehmann, B. N. (1990). Fads, martingales, and market efficiency. *The Quarterly Journal of Economics*, 105(1), 1-28.

Lo, A. (2008). Where Do Alphas Come From?: A New Measure of the Value of Active Investment Management. *Journal of Investment Management*, 6(2), 1-29.

Lo, A. W., & MacKinlay, A. C. (1990). When are contrarian profits due to stock market overreaction?. *Review of Financial studies*, 3(2), 175-205.

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance*, 39(4), 1127-1139.

Securities and Exchange Commission. (2010). Findings regarding the market events of May 6, 2010. *Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues.*

Solman, P. (2012). Fictional Thriller Tackles Dangers of High-Frequency Trading [video]. *PBS Newshour.*

Thal Larsen, P. (2007). Goldman Pays the Price of Being Big. *Financial Times*, August 13.