Examining the Drivers of Optimal Portfolio Construction

A 50 Year Study of the Interaction Between Security Selection and Capital Allocation in US Large Cap Equities, 1968 - 2017

By: Rob Roy

Supervisors:

Prof. Stephen Thomas

Dr. Nicholas Motson

Thesis Submitted for the Degree of

Doctor of Philosophy

Business School, Faculty of Finance

City, University of London

May 11, 2021

Table of Contents

List of Tables 8

List of Figures 9

Acknowledgements 12

Declaration 13

Abstract 14

Introduction 15

1 A Review of the Literature 20

1.1 Introduction . . . 20

1.2 Components of Portfolio Construction . . . 22

1.2.1 Security Selection . . . 22

1.2.2 Capital Allocation . . . 26

1.3 Applications of Active Management . . . 34

1.3.1 Styles . . . 34

1.3.2 Timing . . . 35

1.3.3 Trading . . . 36

1.3.4 Unobserved Actions . . . 37

1.4 Active Share . . . 37

1.5 Methods of Performance Measurement . . . 40

1.5.1 Choosing the Benchmark . . . 40

1.5.2 Choosing the Performance Metric . . . 42

1.6 Performance Attribution . . . 43

1.7 Can Managers Beat the Market? . . . 44

1.7.1 Maybe ‘No’ 45
1.7.2 Maybe ‘Yes’ 46
1.7.3 Performance Inconclusive 47
1.8 Mechanics 48
1.8.1 Data Availability & Quality 48
1.8.2 Methods 49
1.9 Methods of Robustness 50
1.9.1 Construction Tools 50
1.9.2 Performance Verification 51
1.10 Contribution to the Literature 52

2 The Hierarchy of Sharpe Ratios: Observed vs. Significant 55
2.1 Introduction 55
2.2 The Opportunity Sets: Security Selection vs. Capital Allocation 55
2.3 Selection vs. Allocation 58
2.4 Capital Allocation Methods 60
2.4.1 Methods Available But Not Chosen 60
2.4.2 Capital Allocation Methods Chosen For Examination 63
2.4.3 Description of Capital Allocation Methods 64
2.5 Properties of the Returns 72
2.5.1 Returns: 72
2.5.2 Return Distribution Normality: 74
2.5.3 Volatility: 75
2.5.4 Drawdown: 75
2.5.5 Reward vs. Risk: 76
2.5.6 Diversification: 78
2.6 Turnover: 80
2.6.1 Turnover and Transactions Costs 81

2.6.2 Factor Exposures . . . 85
2.7 The Hierarchy of Sharpe Ratios . . . 88
2.7.1 Why Focus on Sharpe Ratio? . . . 88
2.7.2 Hierarchy - Initial Results . . . 91
2.8 Robustness Test - Description . . . 92
2.8.1 Robustness Test - Example Results . . . 95
2.8.2 Robustness Test - Results Summary . . . 97
2.9 Conclusion . . . 99

3 Security Selection & The Contours of Sharpe Ratios 103
3.1 Security Selection . . . 104
3.1.1 Security Selection with Near Perfect Foresight . . . 105
3.2 Positive Selection & Market Capitalization . . . 107
3.3 Simultaneous Security Selection tilt . . . 111
3.3.1 Simultaneous Positive Security Selection . . . 112
3.3.2 Simultaneous Negative Security Selection . . . 113
3.3.3 Security Selection - Full Landscape . . . 114
3.4 Individual Contours - Examples . . . 115
3.4.1 Market Cap vs. Equal Weight . . . 116
3.4.2 Market Cap vs. Minimum Variance 5% . . . 118
3.4.3 Market Cap vs. Maximum Diversification 5% . . . 119
3.5 Landscape Summary . . . 120
3.6 Conclusion . . . 124
3.7 Sharpe Ratio Contours . . . 126
3.7.1 Description of Contours . . . 126
3.8 Market Cap vs. Panel of Allocation Methods . . . 127
3.9 Equal Weight vs. Panel of Allocation Methods . . . 131
3.10 Inverse Volatility vs. Panel of Allocation Methods . . . 134

3.11 Equal Risk vs. Panel of Allocation Methods 137
3.12 Minimum Variance 1% vs. Panel of Allocation Methods 139
3.13 Minimum Variance 5% vs. Panel of Allocation Methods 141
3.14 Maximum Diversification 1% vs. Panel of Allocation Methods 142

4 Low Correlations Are The Dominant Feature Of Optimal Sharpe Ratio Portfolios 144

4.1 Chapter Introduction 144
4.2 Optimal Sharpe Ratio Portfolio - A Model 145
4.2.1 Model Purpose 145
4.2.2 Near-Perfect Foresight 145
4.2.3 OSR Model Definition 147
4.2.4 OSR Portfolio Formation 148
4.2.5 Concentration - Academic & Practical Reference Points 150
4.2.6 OSR Model Design Summary 151
4.3 Performance Attribution of the OSR Portfolio 151
4.3.1 Comparison Portfolios 152
4.3.2 Initial Performance Observations 153
4.3.3 Time Series Isolation 155
4.3.4 Security Selection Crowds Out Diversification Benefits 158
4.4 One Perfect Foresight 160
4.4.1 Introduction 160
4.4.2 Method 161
4.4.3 Observations 161
4.5 Conclusion 165

5 Correlation Selection 168

5.1 Introduction 168

5.2 Refined Method of Selection . . . 169
5.3 Stock Characteristics . . . 171
5.3.1 Mean Stock Returns . . . 171
5.3.2 Mean Stock Volatility . . . 175
5.3.3 Mean Stock Correlations . . . 177
5.3.4 Stock Characteristic Summary . . . 178
5.4 Portfolio Characteristics . . . 179
5.4.1 Portfolio Return . . . 179
5.4.2 Portfolio Volatility . . . 179
5.4.3 Portfolio Sharpe Ratio . . . 181
5.4.4 Portfolio Characteristics Summary . . . 182
5.5 Capital Allocation Methods . . . 182
5.5.1 Capital Allocation Method - Market Capitalization . . . 183
5.6 Sharpe Ratios . . . 185
5.7 Sharpe Ratio Robustness . . . 188
5.8 Conclusion . . . 192
5.9 Time Series Returns of Capital Allocation Methods with Correlation Selection . . . 195
5.10 Sharpe Ratio Robustness Tests . . . 202

6 Research Summary **208**
6.1 Introduction . . . 208
6.2 Hierarchy of Sharpe Ratios . . . 209
6.3 Contours of Sharpe Ratios . . . 210
6.4 Low Correlations Dominate . . . 211
6.5 Correlation Selection . . . 214
6.6 Summary . . . 216
6.7 Future Research . . . 217

7 Appendix 1 **219**

7.1 Introduction . . . 219

7.2 Data . . . 219

7.2.1 Sources . . . 219

7.2.2 Time Frame . . . 220

7.2.3 Exclusions . . . 220

7.2.4 Delisting Returns . . . 221

7.3 Panel of Capital Allocation Methods (Weighting Schemes) . . . 222

7.3.1 Market Capitalization . . . 222

7.3.2 Equal Weight . . . 223

7.3.3 Inverse Volatility . . . 224

7.3.4 Equal Risk Contribution . . . 225

7.3.5 Minimum Variance . . . 227

7.3.6 Maximum Diversification . . . 228

7.4 Statistical Measures . . . 230

7.5 Comparison of Capital Allocation Methods . . . 236

7.5.1 Statistical Characteristics . . . 236

7.5.2 Factor Exposures . . . 245

7.6 Computational Intensity . . . 248

References **250**

List of Tables

1 **Summary statistics for various capital allocation methods, 1968 - 2017** . . . 73

2 **Marginal Trading Cost Analysis** . . . 84

3 **Fama-French Three Factor Model, 1968 - 2017** . . . 86

4 **Sharpe Ratio Differences** . . . 93

5 **Significance Testing of Monthly Sharpe Ratio Differences** . . . 97

6 **Monthly Sharpe Ratio Comparisons With Security Selection: Market Cap method** . . . 109

7 **Changes in the significance test of monthly Sharpe ratio differences in the presence of positive security selection tilt** . . . 121

8 **Changes in the significance test of monthly Sharpe ratio differences in the presence of negative security selection tilt** . . . 122

9 **Time Series Summary - Market Capitalization** . . . 185

10 **Time Series Summary - Equal Weight** . . . 195

11 **Time Series Summary - Inverse Volatility** . . . 196

12 **Time Series Summary - Equal Risk** . . . 197

13 **Time Series Summary - Minimum Variance, 1%** . . . 198

14 **Time Series Summary - Minimum Variance, 5%** . . . 199

15 **Time Series Summary - Maximum Diversification, 1%** . . . 200

16 **Time Series Summary - Maximum Diversification, 5%** . . . 201

17 **Summary statistics for various capital allocation methods, 1968 - 2017** . . . 237

18 **Fama-French Three Factor Model, 1968 - 2017** . . . 246

List of Figures

1 **Annual Returns - Top 1,000 Largest Stocks** 56
2 **Mean Annual Correlation - Top 1,000 Largest Stocks** 57
3 **Hierarchy of Sharpe Ratios** 91
4 **Statistical Significance of Sharpe Ratio Differences** 95
5 **Upward tilted market capitalization weighting method** 107
6 **Sharpe ratios with Positive Security Selection tilt** 112
7 **Sharpe ratios with Negative Security Selection tilt** 114
8 **Sharpe Ratios with Positive and Negative Security Selection Tilt** 115
9 **Sharpe ratio contour: MC vs. EW** 116
10 **Sharpe ratio contour: MC vs. MV5%** 118
11 **Sharpe ratio contour: MC vs. MD5%** 119
12 **Sharpe ratio contour: MC vs. EW** 127
13 **Sharpe ratio contour: MC vs. IV** 127
14 **Sharpe ratio contour: MC vs. ER** 128
15 **Sharpe ratio contour: MC vs. MV1** 128
16 **Sharpe ratio contour: MC vs. MV5** 129
17 **Sharpe ratio contour: MC vs. MD1** 129
18 **Sharpe ratio contour: MC vs. MD5** 130
19 **Sharpe ratio contour: EW vs. IV** 131
20 **Sharpe ratio contour: EW vs. ER** 131
21 **Sharpe ratio contour: EW vs. MV1** 132
22 **Sharpe ratio contour: EW vs. MV5** 132
23 **Sharpe ratio contour: EW vs. MD1** 133
24 **Sharpe ratio contour: EW vs. MD5** 133
25 **Sharpe ratio contour: IV vs. ER** 134
26 **Sharpe ratio contour: IV vs. MV1** 134

27 **Sharpe ratio contour: IV vs. MV5** 135
28 **Sharpe ratio contour: IV vs. MD1** 135
29 **Sharpe ratio contour: IV vs. MD5** 136
30 **Sharpe ratio contour: ER vs. MV1** 137
31 **Sharpe ratio contour: ER vs. MV5** 137
32 **Sharpe ratio contour: ER vs. MD1** 138
33 **Sharpe ratio contour: ER vs. MD5** 138
34 **Sharpe ratio contour: MV1 vs. MV5** 139
35 **Sharpe ratio contour: MV1 vs. MD1** 139
36 **Sharpe ratio contour: MV1 vs. MD5** 140
37 **Sharpe ratio contour: MV5 vs. MD1** 141
38 **Sharpe ratio contour: MV5 vs. MD5** 141
39 **Sharpe ratio contour: MD1 vs. MD5** 142
40 **OSR Portfolio: Number of Stocks Selected** 149
41 **OSR Performance Comparison** 153
42 **Average Security Performance Comparison** 155
43 **Average Security Volatility Comparison** 156
44 **Average Pairwise Correlation Comparison** 157
45 **Optimizing the 20 Best Performing Stocks** 159
46 **Selection and Weighting of Optimize 20 Best Performing Stocks** . 160
47 **The Landscape of Perfect Foresight by Parameter** 162
48 **Portfolios Formed With One Perfect Foresight** 164
49 **Change in Mean Stock Return - Tilt to Lower Correlation** 172
50 **Change in Mean Stock Return - Tilt to Higher Correlation** 173
51 **Correlation Selection: Mean Stock Return Comparison** 174
52 **Change in Mean Stock Volatility - Tilt to Lower Correlation** . . . 175
53 **Change in Mean Stock Volatility - Tilt to Higher Correlation** . . . 176

54 **Correlation Selection: Mean Stock Volatility Comparison** 177
55 **Change in Mean Stock Correlation** 178
56 **Change in Portfolio Volatility** 180
57 **Change in Portfolio Sharpe Ratio** 181
58 **Time Series of Returns: Market Capitalization** 185
59 **Panel of Allocation Method Sharpe Ratios Across Correlation Selection** 186
60 **Sharpe Ratio Robustness Test: Equal Weight** 189
61 **Sharpe Ratio Robustness Test: Minimum Variance 1%** 190
62 **Time Series of Returns: Equal Weight** 195
63 **Time Series of Returns: Inverse Volatility** 196
64 **Time Series of Returns: Equal Risk** 197
65 **Time Series of Returns: Minimum Variance, 1%** 198
66 **Time Series of Returns: Minimum Variance, 5%** 199
67 **Time Series of Returns: Maximum Diversification, 1%** 200
68 **Time Series of Returns: Maximum Diversification, 5%** 201
69 **Sharpe Ratio Robustness Test: Market Capitalization** 202
70 **Sharpe Ratio Robustness Test: Inverse Volatility** 203
71 **Sharpe Ratio Robustness Test: Equal Risk** 203
72 **Sharpe Ratio Robustness Test: Minimum Variance 5%** 204
73 **Sharpe Ratio Robustness Test: Maximum Diversification 1%** . . . 204
74 **Sharpe Ratio Robustness Test: Maximum Diversification 5%** . . . 205
75 **Correlation Selection Time Series Summary Statistics** 206
76 **Number of stocks available in the CRSP database** 220

Acknowledgements

While growing up, my family instilled curiosity in me and helped launch me into a self-sustaining life of learning. They are due many thanks for this gift. I must thank everyone who kindly suffered my questions over the last 40+ years. This list includes professors, investment managers, colleagues, friends and family. Their graciousness allowed me to grow and indulge a sense of wonder for this world of portfolio management. In an industry often focused on the acquisition of wealth, there have been several people who seek to understand the world more fully, and pushed ego aside to seek the best answers. A higher level of thanks is for those who went beyond indulging my curiosity and pushed back at my questions, uncomfortably leading me to examine my own thinking even more deeply. I especially appreciate my colleagues at AdventHealth who have believed in me enough to support this project, push me to be better each day, and provide a shining example of Christian service.

Thank you also to my two wonderful daughters, Alix and Coco. I hope that seeing an old man using and developing the talents given to me by God will be an inspiration to you throughout your life. But the highest level of thanks goes to my wife and best friend who understands me and appreciates my curiosity and drive to improve as my act of thankfulness to the world around us.

Declaration

I grant powers of discretion to the University Librarian to allow this thesis to be copied in whole or in part without further reference to me. This permission covers only single copies made for study purposes, subject to normal conditions of acknowledgement.

Abstract

Markowitz advanced the theory of portfolio construction, dividing the work into two stages of security selection and capital allocation. He did not address security selection, but established a new method to allocate capital. Since then, much has been written regarding each stage while the investment industry grew and its performance was examined through a variety of lenses. This thesis contributes to the literature by examining the interaction between these two stages of portfolio construction. Each stage impacts performance, but an examination of the interaction is missing from the literature. A deeper understanding of the interaction provides opportunity for practitioners to improve portfolios and gives a base for future research to refine the understanding.

The research begins by holding 1,000 US large capitalization stocks constant while applying a panel of capital allocation methods. Despite starting a fixed set of securities, differences in performance and risk characteristics are documented. Markowitz identified the objective superiority of risk adjusted returns over just returns, so my work applies the Sharpe ratio as the measurement unit. Differences in Sharpe ratios across the panel of allocation methods are tested for robustness using a bootstrap test.

With a hierarchy of Sharpe ratios established, the next step varies the security selection and observes the change in Sharpe ratios. Security selection is implemented with one year perfect foresight which is a limit condition for the potential of stock picking. When applying good and bad security selection it is observed that the hierarchy of Sharpe ratios is unstable in the presence of security selection. The bootstrap robustness test shows the hierarchy can invert with statistical significance. This is the first step in understanding the dependence of the optimal capital allocation method in the presence of security selection skill.

Toward the goal of optimal Sharpe ratio portfolios, I examine portfolios constructed with perfect foresight into return, volatility, and correlations. Optimized portfolios are dominated by low correlation securities, not by high returning securities. Conversely, by first selecting the highest returning securities possible, you crowd out the optimal Sharpe solution. Reward curves are built for selecting securities based on returns, volatility, and correlation. The shape of the correlation reward curve is like the low volatility anomaly.

Because low correlation securities dominate optimized Sharpe ratio portfolios, perfect correlation selection is applied. Stock level and portfolio level characteristics are documented in the cross section of years. Then the method is applied to the panel of allocation methods to create return time series and Sharpe ratios. The bootstrap robustness test applied to the Sharpe ratios of the allocation methods shows that even in the absence of perfect correlation selection, nearly any level of success achieved in correlation selection creates robust improvements in the Sharpe ratio.

Introduction

Over the last 30+ years, I have been responsible for both internal management of capital and for allocating capital to external money managers. Inevitably, the presentations from external managers turn into ‘story-time’ and I am regaled with anecdotes about the most interesting companies that they have ‘discovered’ and how this skill has led to their superior returns. Usually I don’t contest their stock picking skills, but eventually I ask them how they allocate capital across their selected securities. The most common answer given is that they ’conviction weight’ their holdings.

Conviction weighting is simply a discretionary method used to allocate more money into the stocks that they feel most strongly about. The academic literature for equity portfolios is well developed in regard to systematic capital allocation methods outperforming relative to basic market capitalization weights. Researchers such as (Brinson, Hood, and Beebower 1995) and others have examined the question of selection versus allocation in the setting of multiple asset classes. The conclusions of these studies indicate that the allocation decision is significantly more important than selection. If this were true when focused on an equity portfolio, it would be curious why managers spend so much time on the task of security selection. This gap between the academic literature and the manager stories incited my curiosity. What are the connections between security selection skills and capital allocation methods? What drives optimal portfolio formation? What implications does this have for active equity managers?

In the research presented here, I have pursued these questions theoretically by examining the intersection between a universe of securities with and a panel of allocation methods

Through my research I have found that within the single asset class of equi-

ties, security selection can be helpful but is not the dominant path toward improvement. In addition, I find that the optimal choice of an allocation method is dependent on proficiency in security selection.

To begin, I start back at the beginning of portfolio construction with Harry Markowitz who projected the complexity of portfolio construction on to the two tasks of security selection and capital allocation, focusing his work on the latter.

> *"The process of selecting a portfolio may be divided into two stages. The first stage starts with observation and experience and ends with beliefs about the future performance of available securities. The second stage starts with the relevant beliefs about future performances and ends with the choice of portfolio." (Markowitz 1952)*

Markowitz's research focused on the second stage of portfolio construction: capital allocation. He identified the benefits of incorporating the relationships (correlation) between securities to optimize the risk adjusted return of a portfolio. Separate strains of literature and work have focused on each of the two stages, but the connection between each stage has not been extensively explored. My research connects these two stages and focuses on the relationship between security selection and capital allocation in pursuit of designing more optimal portfolios.

This research is important to both investment managers as well as investors who allocate their capital to professional managers. For investment managers, understanding this relationship more clearly can lead to a more effective focus of resources which in turn may lead to improving the outcomes for their investors. For investors who allocate, the methods identified in this research allow for a more detailed analysis of manager returns.

Chapter 1 is a survey of the existing literature related to the themes of this research project.

Chapter 2 begins by focusing on the potential value available from well known and commercially available allocation methods. Across a 50 year history of US large cap stocks, a hierarchy of Sharpe ratios emerges suggesting that value can be created by following allocation methods which differ from the more common method of market capitalization weighting. But what appears evident in the time series is not always statistically significant. Across the full matrix of results, an advanced method for checking the significance of the Sharpe ratio differences is applied which shows that 'all that glitters is not gold'.

Chapter 3 extends the research by introducing security selection into the analysis. The application of 'near perfect foresight' is applied to the universe of securities. The first application of this method is to show the minimum amount of security selection expertise required, in the limit, to sufficiently improve the market cap weighting when compared against the panel of alternative allocation methods. The tool is then extended further to each of the allocation methods to show that the conclusions of statistical significance related to the hierarchy of Sharpe ratios established in Chapter 2 is not invariant to security selection. The important conclusion is that the optimal choice of an allocation method is highly dependent on the level and persistence of security selection skill.

Chapter 4 finds that portfolios with the highest possible Sharpe ratios are dominated by low correlation securities, not by the highest returning securities. In fact, by building a subset of stocks with the highest returns each year, you crowd out the potential for maximizing the Sharpe ratio. The reward curve of Sharpe ratio, by selecting securities based on performance is quite steep, but bad selections can offset good selections. By contrast, the reward curve for picking stocks based on volatility is positive across all four quartiles and is confirmed when combining the stocks into a portfolio. Using what we see from the creation of optimal Sharpe ratio portfolios, I examine the reward curve for portfolios where security selection is based on pair-wise correlations. The shape of the correlation based reward curve is like that of the volatility reward curve, but is even steeper when selecting the lowest quartile of pair-wise correlations.

Chapter 5 applies the methods from Chapter 3 and examines the impact of applying perfect foresight to selecting securities based simply on pair-wise correlations. The effect of this method is examined first in the cross section of years and then in a time series approach. Bootstrap robustness tests are applied to show the significance of this method at nearly any level of success.

Each of the chapters highlights the importance of the interaction between security selection and capital allocation when forming a portfolio. For the researcher, these papers can point the way toward developing better methods to integrate volatility and correlation forecasts. For the practitioner, these papers show the importance of understanding and connecting these two stages of portfolio formation while highlighting the smaller role that security return forecasting should play in their process.

Chapter 1

A Review of the Literature

1 A Review of the Literature

1.1 Introduction

As I will review in this section, there exists a growing body of literature which examines performance of portfolios using a variety of lenses. Much of the work is empirical and tries to explain the shadows created by managers who have attempted to build optimal portfolios through time. Some of the literature is theoretical and focused on either security selection or capital allocation.

This thesis expands on the existing literature by examining the interaction effect between security selection and capital allocation while highlighting where significant improvements are possible and refuting the practitioner's dominant focus on return forecasting.

My literature review begins with the work of (DeMiguel, Garlappi, and Uppal 2009), who points us toward a fourth century rabbi who espoused portfolio formation with an equal allocation methodology. It is also noted that portfolio construction is found in the Babylonian Talmud[1] where an equally weighted, three asset portfolio is espoused: *a third in land, a third in merchandise, and a third ready to hand.* However, not much else is written on this topic until (Markowitz 1952) is published.

In the first lines Markowitz states:

> *"The process of selecting a portfolio may be divided into two stages. The first stage starts with observation and experience and ends with beliefs about the future performance of available securities. The second stage starts with the relevant beliefs about future performances and ends with the choice of*

[1] `https://www.jewishvirtuallibrary.org/babylonian-talmud-full-text`

portfolio." (Markowitz 1952)

With this introduction, the search for portfolio improvements begins and continues to this day.

Although Markowitz laid the foundations of portfolio selection, scholars since have not been in unison on which methods to use, the practicality of implementing the methods, the approach to measuring performance, or the overall success of the investment management industry. The breadth of literature makes clear there exist multiple methods for constructing portfolios and for measuring the outcome of these applications.

The literature will be reviewed within the following framework:

- Components of Portfolio Construction
  - Security Selection
  - Capital Allocation
- Applications of Active Management
  - Styles
  - Timing
  - Trading
  - Unobserved Actions
- Active Share
- Methods of Performance Measurement
  - Measuring Activity
  - Choosing the Benchmark

  - Choosing the Performance Metric(s)

- Can Managers Beat the Market?
  - Maybe 'No'
  - Maybe 'Yes'

- Mechanics
  - Data Availability & Quality
  - Calculation Methods

- Robustness
  - Construction Tools
  - Performance Verification

1.2 Components of Portfolio Construction

Markowitz set the task of portfolio selection into two main stages: security selection and capital allocation. While the literature related to each component is reviewed below, these are not entirely separate and distinct steps. The authors of (Amenc, Goltz, and Lodh 2012) point out that methods within each step of the process have various advantages and drawbacks, they are not competing elements but can be used successfully in combination to reach the investor's objective. This is a key observation which my research expands upon.

1.2.1 Security Selection

> *"The first stage starts with observation and experience and ends with beliefs about the future performance of available securities." (Markowitz 1952)*

While Markowitz focused his paper on the capital allocation question he did note that a best method for security selection was likely to be the combination of statistical techniques and the 'judgment of practical men'. The combination of these two elements has been greatly expanded in the literature since that time.

Portfolios are simply collections of individual securities selected from the array of all securities available across the globe. While the first step is security selection, often the opportunity set is narrowed by a stated objective of the fund or by an investment policy rule. Criteria to narrow an opportunity set may include country/geography, asset class, market capitalization, or factor exposure. An example of a restricted opportunity set is US Equity Large Capitalization Value. The authors of (Clarke, Silva, and Sapra 2011) caution however, that investors should be aware of the inefficiencies that can be created by each of the constraints and the interaction effect of multiple constraints.

Within the literature related to security selection, there are documented many systematic yet fundamentally based strategies. In the early part of this literature, (Treynor and Black 1973) picked up this thread from Markowitz and applied a level of mathematical rigor to security selection. Today there exists an enormous literature from the practitioner's side which suggests a variety of methods for picking better stocks. As computational capabilities have increased, this literature is now flavored with elements of artificial intelligence and other advanced data science methods. This dissertation is not focused on the method to pick to stocks correctly, but instead focuses on the interdependence of this skill with the optimal choice for capital allocation method and the drivers of optimal Sharpe ratio portfolios.

1.2.1.1 Discretionary Security Selection The absence of a systematic approach to security selection is by definition discretionary. An investment manager may be following a heuristic, but the lack of any systematic method gives rise to a category which can not be analyzed directly. Despite not following an established set of decision rules, literature related to characteristics of the equity market may be guiding the success of managers who follow a discretionary approach to security selection.

An example of these equity market characteristics is (Ang et al. 2006) who examine the cross section of volatility and expected returns to find that stocks with high idiosyncratic volatility have abysmally low average returns. This follows on from the earlier work of (Fama and French 1992) who began to show that a stock's exposure to common fundamental factors such as size, and book to market, could help to explain much of the difference in returns.

Idiosyncratic risk arises in security selection, but a deeper understanding of this component may impact returns. This insight was offered by (Goyal and Santa-Clara 2003) who find a significant positive relation between average stock variance and the return of the market. The authors find that variance of the stock market by itself has no forecasting power over the market return. However, idiosyncratic risk explains most of the variation of average stock risk through time and it is idiosyncratic risk that drives the forecastability of the stock market. This finding is at odds with the theoretical literature which posits that only systematic risk should be priced in the markets as investors can easily diversify away idiosyncratic risk.

According to (Sebastian and Attaluri 2016), exposure to factors seems to have explained a large degree of superior performance in strategies through time. However, it is the discretionary skill of adaptation that may have led to persistence in performance.

There is evidence that a discretionary approach to selection may add value. In the work of (Baker et al. 2010), the authors examine the performance of stocks traded by mutual funds around the announcement of earnings. New purchases tend to outperform similar stocks and new sales tend to under-perform similar stocks. This skill however seems to be waning since the passage of Rule-FD from the SEC which mandated fair and even disclosure of information to investors.

1.2.1.2 Systematic Security Selection Following a clear set of rules defines the category of systematic security selection. These methods are highly varied, but as they have fully disclosed rules, they can be analyzed and their properties examined.

Departing from the common starting point of security selection based on capitalization, (Arnott, Hsu, and Moore 2007) introduce the concept of fundamental indexation which begins to narrow the set of securities according to their fundamental characteristics such as book value, cash flow, revenue, sales, and dividends. The soundbite logic is provided by Benjamin Graham who said:

> *"In the short run, the market is a voting machine, but in the long run, it is a weighing machine."*

The success of this approach appears to be another empirical violation of the simple CAPM model. This work was greatly expanded by (Clare, Motson, and Thomas 2013).

The explanation of success through fundamental characteristics of stocks in the Arnott paper is refuted by (Perold 2007) who find that fundamental indexing is just a form of quantitative value investing and if value stocks are systematically mispriced, this is the source of the return and not the headwind of capitalization weighting.

Depending on the investment beliefs of the manager, other systematic strategies may be designed. Examples may include focuses on growth stocks or deep value stocks. While the literature documents many of these approaches, it is an open topic on whether security selection matters at all. A question examined in this thesis asks to what degree security selection impacts overall performance relative to the impact of capital allocation and how the two stages interact to impact the risk-return efficiency of a portfolio.

1.2.2 Capital Allocation

> *"The second stage starts with the relevant beliefs about future performances and ends with the choice of portfolio." (Markowitz 1952)*

After securities are selected for inclusion, the next stage in the process is to allocate capital across the securities. The general objective is to find the best portfolio for the next period with an optimal trade-off between return and risk. Without imposed limits, the possible allocation weight combinations are infinite. Often however, there are limits on characteristics such as: short vs. long positions, minimum and maximum position size, leverage constraints, etc. Within the existing constraints, managers must find a way to allocate their capital across the selected securities.

Active equity managers appear to spend an excess of resources on analysis for security selection relative to the effort expended on considering the method of capital allocation. This stands in contrast with (Clare, Motson, and Thomas 2013), (DeMiguel, Garlappi, and Uppal 2009) and others who show that by simply abandoning a market capitalization weighting scheme and following an alternative allocation methodology, superior results have been possible over a long time period. The form of capital allocation most available to index investors today is capitalization weighting. In this approach, the companies with the largest market value are given the greatest weight in the portfolio. This is not necessarily the most efficient way to allocate capital, but it does have characteristics of high liquidity, low transaction volumes, and low costs that make it attractive. Deviations from this method exist in both discretionary and systematic formats. Capital allocation is in the domain of the active portfolio manager.

Despite the predominance of products based on capitalization weighted indices, it is commonly attacked in the academic and practitioner literature. One example, (Hsu 2006) shows that the sub-optimality arises because cap-weighting tends to overweight stocks whose prices are high relative to their fundamentals and underweight stocks whose prices are low relative to their fundamentals. In essence, the cap-weighted index is 'anti-value'. A detailed mathematical proof shows that the cost of sub-optimality is equal to the square of the noise of the stock prices. Another example of this decomposition is shown in (Martellini 2008) who shows that the efficiency of cap-weighted indices can only be warranted at the cost of heroic assumptions.

Two papers in particular were helpful in the overview of these various approaches. The work of (Kolm, Tütüncü, and Fabozzi 2014) and (Cremers et al. 2018) surveyed the literature over the last several decades and illuminated the expansion of new allocation methods.

1.2.2.1 Discretionary Capital Allocation This form of allocation is defined simply as the absence of a systematic approach. Any form of human judgment applied would be considered discretionary. Through many manager interviews I have heard them describe their portfolio as 'conviction weighted'. This particular form of allocation gives more capital to a managers 'best ideas' which may change through time. Often these methods rely on the experience of the individual manager and their interpretation of current market conditions. In addition, the conviction weighting is being driven by the first stage of security selection and not by the optimal combination of the relationships (correlation) between securities.

Most active managers understand that correlations are an important component of total portfolio volatility, but my research will show that correlations are the primary potential driver of optimal Sharpe ratio portfolio and should also be incorporated during the security selection stage.

The discretionary form of allocation is not studied within this thesis as the methods are not disclosed, may be influenced by behavioral issues, and are hard to categorize. Anecdotally, these discretionary approaches appear to be the most predominant form of allocation method followed by active US equity mutual funds.

1.2.2.2 Systematic Capital Allocation Beyond discretionary allocation methods, there is a growing set of systematic capital allocation methods which are well described and examined in the literature. The work of (Clare, Motson, and Thomas 2013) sorted these methods into sensible categories of heuristic, optimization, and random methods.

Heuristic: Heuristic forms of allocation follow simple rules that can be calculated in closed form. Examples of this category are equal weighting, inverse volatility, and diversity weights, among others.

Optimization: Optimization forms of allocation cannot be calculated in closed form and must be solved by numerical routine. Examples of this type are minimum variance, maximum diversification, and risk efficient weights, among others.

Random: An unlimited number of random methods could be applied to the allocation decision, however (Clare, Motson, and Thomas 2013) introduce the monkey portfolio and the Scrabble weight as examples in this category.

Starting at the beginning, (Markowitz 1952) explained the use of a mean-variance optimization method, and despite its relative simplicity, it is still not widely used in practice. (Michaud 2014) points out that optimizers make conceptual demands on portfolio managers and may move counter to politics which often exist within organizations. In addition, these methods require knowledge of statistical methods and are not necessarily intuitive. This may help to explain the predominance of conviction weighting by active US equity managers.

(Clare, Motson, and Thomas 2013) use a panel of systematic allocation methods to observe general performance characteristics when applied to large capitalization US equity portfolios from 1968 through 2011. All of the alternative capital allocation methods (including random weighting schemes) would have produced a better risk-adjusted performance than a capitalization weighted approach, yet there isn't a single method that dominates all other methods in all time periods. Each method has a set of assumptions under which the method proves to be optimal.

The authors follow the established methodologies in the literature where they existed. For Diversity Weights, they follow the work of (Fernholz, Garvy, and Hannon 1998), while for equal risk contribution they reference the work of (Maillard, Roncalli, and Teiletche 2008). Finally, for maximum diversification method, the work of (Choueifaty and Coignard 2008) is cited.

Departing from the methods already established in the literature, (Clare, Motson, and Thomas 2013) test a random selection method which they call the "monkey" portfolio by randomly drawing with replacement from a pool of available securities and giving each draw 0.1% weight until all capital is allocated. The authors create ten million runs of this type and examine the superiority of these random portfolios over the cap-weighted approach.

(Amenc, Goltz, and Lodh 2016) takes issue with the superiority of random "monkey" portfolios and shows that the strategies employed by commercially available approaches, termed 'Smart Beta' products are not 'monkey business'. The authors also criticize the work of a related approach, (Arnott et al. 2013) which looked at inverting the logic of various allocation methods. The criticism exists because the authors did not invert the stock selection step in their analysis. Once this is added to the analysis, the inverted strategies do worse.

Diversification is said to be the 'only free lunch on Wall Street'. Reducing the risk of a portfolio without necessarily diminishing the returns is an improvement by Markowitz's definition. But this concept of diversification previously lacked specifics. A formal measure of diversification is introduced by (Choueifaty and Coignard 2008) which they called the Diversification Ratio. This ratio is defined as the weighted sum of the risks of each security divided by the risk of the combination of the weights. Formally, this is:

$$Diversification\ Ratio = \frac{\sum_{i=1}^{n} w_i \sigma_i}{\sigma_p}$$

where w_i is the weight of stock i, σ_i is the standard deviation of stock i, and σ_p is the standard deviation of the portfolio combination of all stocks.

The full set of portfolio invariance properties of the diversification ratio are laid out in (Choueifaty, Froidure, and Reynier 2011).

(Chow et al. 2011) conducted observational work on a similar set of methods, but with an extension across US and global equity markets. Their work found that any of these systematic allocation methods can be mimicked by combinations of the other strategies which spanned the market, value, and size factors. With this finding and acknowledging that there were already too many methods available to follow, they suggest that the analysis of implementation costs are a better method of evaluation than are back-test returns. In sample, returns comparisons can show differences, but with the addition of robustness checks, ex ante there can be no conclusion of the return dominance of a single systematic method.

In Chapter 2 of this thesis, the same robustness methodology is applied to the differences in Sharpe ratios with a contrasting conclusion. This implies that systematic capital allocation methods may not create significant differentials in return but can create significant improvements in risk-return efficiency.

The most simplistic of the systematic methods, equal weighting, is put forth in the very rigorous (DeMiguel, Garlappi, and Uppal 2009) as the benchmark by which to measure all other approaches, both systematic and discretionary. In this simple method, capital is allocated equally among all selected securities. The authors examine 14 capital allocation methods against seven different data sets and make observations based on returns, Sharpe ratios, and certainty equivalents.

Nearly all systematic allocation methods require estimating parameter inputs such as volatility and correlation. A major contribution of this work is showing an analytical method for correctly estimating the length of period required for estimating the parameters of the simple mean-variance method which would out-perform an equally weighted method. The required period would be excessively longer than any person's career or lifetime. For a 50 asset portfolio, the period needed would approach 6,000 months, which is 500 years. The authors conclude that equal weighting is a superior approach except when the number of assets is very large and an optimization method may find a more diversified combination.

Beyond simply describing methods and examining back-tests of systematic capital allocation methods, several authors including (Maillard, Roncalli, and Teiletche 2008) and (Choueifaty, Froidure, and Reynier 2011) explored the mathematical properties of these approaches and showed the similarities, differences, and highlighted the specific assumptions under which each approach is optimal.

And yet with all of the approaches explained and analyzed, it is still incredibly hard for investors to select and access these approaches (Amenc, Goltz, and Lodh 2012). In this work they show a wide variety of commercially available offerings and the lack of transparency given, the layers of complexity arising from the interaction of multiple decisions, and the inability to examine the track record shown. Within this important work, they refute that the out-performance of the systematic capital allocation approaches is simply due to exposure to small cap stocks or other known factors. By restricting the security set, significantly superior performance still exists. This thesis will extend this work by examining the interaction between security selection and capital allocation. Robustness checks are applied to the results.

One of the more obscure allocation methods was described by (Bera and Park 2008) who introduce diversification using the maximum entropy principle. The method attempts to account for the non-normality of financial data and the effect this has during the specification of the covariance matrix. The approach also attempts to overcome the well known issues with mean-variance optimization, which include the low diversification and concentration issues that arise as a function of misspecified covariance matrices.

The maximum entropy approach begins with a realization that the maximum likelihood approach is inconsistent due to not knowing the true density function. It then attempts to "extract useful information about the unknown density from a given data by imposing some well-defined moment functions in analyzing financial time-series data." By doing this, the authors show that model misspecification can be reduced. This new density is the maximum entropy density which is used in the Engle/Bollerslev style ARCH framework.

The authors also compare their approach to the Bayes Stein approach which uses estimator shrinkage. This work is significant because it is but a sample of works which point out the shortcomings of using a point estimate approach to creating the inputs to optimizations. By thinking of the parameters of return, volatility, and correlation as probability distributions, rather than specific values, their approach accounts for the inherent uncertainty in these values and directly shrinks the optimized weights toward a reasonable prior, such as market cap weighted or equal weighted.

Attempting to refute the source of value emanating from a systematic allocation method, in (Arnott et al. 2013), the authors invert the logic of several strategies. If the original logic produced superior returns over cap-weighted indices, then inverting the rules should produce worse results, but this was not the case, and they generally performed better than the original strategies. The source of the superior performance is claimed to be created from factor exposure to small cap and value tilts.

In a multi-asset class setting, (Chaves et al. 2011) tests the properties of the risk parity approach and the conditions under which this method would excel. While the approach is sensitive to asset class selection, returns do not need to be estimated. However, the authors fail to address the sensitivity issues related to estimating the parameters of the covariance matrix.

In another paper comparing systematic approaches, (Clarke, Silva, and Thorley 2012) examine the computational requirements of three major approaches. The restrictions of computer processing speed seem to have been alleviated since the publication of this article, but the authors create a framework to logically explain why any given asset is included in the risk based portfolios and show the reasoning for the weight computed.

The comparison of approaches is simplified by (Carvalho, Lu, and Moulin 2012) where the authors begin with an equal weight portfolio and show how each approach is simply a nested case of the preceding strategies. The net exposure differences between nested methods are based on simple beta and alpha tilts. This is a good explanation for less mathematically inclined practitioners.

An outside perspective is offered by (Fabozzi, Huang, and Zhou 2010) who publish in the Annals of Operations Research. While the omnipotence of the mean-variance approach exists in theory, the practice is not as ubiquitous as expected. Referencing Robert Merton who said that the challenge is to put “the rich set of tools” into practice but concluded that “I see this as a tough engineering problem, not one of new science.” The operations research angle means that the authors take the parameter estimates as given and examine the robustness of the methods applying these estimates. This is a very detailed mathematical paper and exposes many of the strengths and weaknesses of the various systematic allocation methods.

The collection of systematic methods continues to grow and with it the complexities of interaction terms from the layers of decisions. While it seems evident that efficiency is gained from this class of allocation method, the existing literature does not examine the relative importance of these methods versus security selection. The analytical work in this thesis will cover this gap in the literature.

1.3 Applications of Active Management

The literature reviewed to this point outlines the major methods for constructing portfolios. The next strain of literature examines empirical applications of these tools. The pursuit of superior performance requires the application of a combination of active management tools which will create different, and hopefully superior, performance.

1.3.1 Styles

Active management of funds may be implemented using particular styles of investing. Rather than analyzing the full universe of potential securities, managers may limit the scope of their investment to companies that display some combination of characteristics. The particular characteristics chosen form the manager’s style.

The work of (Fama and French 1992) helped to explain the dispersion in stock returns that was not being captured by the simple asset pricing models of (Sharpe 1964) or (Lintner 1965). It was observed that a regression of stock returns on the returns of the general market (β) was inadequate to explain the full set of returns. Factors such as the market capitalization or book to price were able to more fully explain the cross section of stock returns. In a widely cited work, (Carhart 1997) added to the factor model work with the inclusion of momentum.

Investment beliefs have been formed from the historical observations of these style factors and are used today in the management of funds with particular objectives. A recent book by (Ang 2014) captures the practitioner view of these style investing processes and how they are implemented.

1.3.2 Timing

Well before much of the literature referenced herein, (Fama 1972) examined methods possible for deciphering the effects of 'selectivity' from the effects of 'timing'. In this regard, timing is the ability to predict general market price movements and adjust a portfolio to benefit from this activity. Later, (Cremers and Petajisto 2009) also cite the ability for managers to adjust their holdings and benefit from timing the market.

Identifying managers who can beat the market through timing is a direct challenge to the efficient market hypothesis. However, using daily data, (Bollen and Busse 2001) show that approximately 40% of managers in their study were able to add value through market timing. This is a direct contradiction of the previous work of (Treynor and Mazuy 1966) who show no evidence of managers ability to time the market correctly.

Finally, (Jiang, Yao, and Yu 2007) use portfolio holdings to examine whether timing skill is applied successfully. The authors find that on average managers do have positive timing ability at the one month and twelve month horizons. Additional statistical robustness checks are added over previous literature. The funds with the best timing ability tend to have greater industry concentrations and are tilted toward small cap stocks. Funds that become too big can't trade enough stock to impact their overall returns.

It is the conclusion of (Huang, Sialm, and Zhang 2011) that funds that engage in timing, by shifting their risk dynamically through time, fare worse than funds that maintain stable risk profiles.

1.3.3 Trading

The effect of active trading was examined by (Wermers, Chen, and Jegadeesh 2000). The authors showed that preference for liquidity may be hurting fund manager returns and that trade performance is positive for growth funds but insignificant for income oriented funds. High turnover funds seem to capitalize on their stock selection abilities by trading and timing frequently. Stocks that are widely held by mutual funds do not outperform the general population of stocks but more heavily traded stocks can add positive returns.

The authors of (Pastor, Stambaugh, and Taylor 2017) find that the time series relation between turnover and performance is stronger than the cross sectional relation. In addition, they find that the turnover relation to performance is stronger for funds trading less liquid stocks while the level of turnover is correlated across funds. Profit opportunities vary over time and some funds exhibit the ability to identify and exploit these opportunities. While the active management industry may not provide superior net returns to its investor (consistent with both theory and evidence), it creates a valuable externality in that the combined trading of active funds helps to correct prices and thereby enables more efficient capital allocation.

1.3.4 Unobserved Actions

Due to the difference in timing between quarterly portfolio holdings disclosure and daily net asset value (NAV) reporting, not all performance can be attributed because some actions are unobserved. In the work of (Kacperczyk, Sialm, and Zheng 2008), the authors show that the average gap in performance is near zero which implies that the aggregate result of interim manager actions creates virtually no value but is just enough to offset transaction costs. However, past return gaps tend to persist so funds that add value through unobserved actions continue to add value while those that destroy value continue as well. It is impossible to disentangle what actions are being taken during the gap of time and the return gap that is created.

1.4 Active Share

In an attempt to create superior performance, the applications of active management strategies create portfolios that are different than their benchmarks. In 2009, a new metric called **Active Share** (AS) was introduced by (Cremers and Petajisto 2009) and has now gained widespread commercial acceptance with firms such as Morningstar and many financial advisors. The AS metric was intended to measure how 'different' a given portfolio is from its comparison benchmark and to create a link between the level of 'active management' and the potential for improved returns.

The stages of security selection and capital allocation are combined into a single measure that seems to be intuitive. The measure is based on current holdings of the portfolio rather than being based on historical return differences that are used to calculate tracking error.

The mathematical representation of Active Share is:

$$ActiveShare = \frac{1}{2}\sum_{i=1}^{N}|w_{fund,i} - w_{index,i}|$$

where $w_{fund,i}$ and $w_{index,i}$ are the portfolio weights of asset i in the fund and in the index, and the sum is taken over the universe of all assets included in the fund and the index.

Tracking error is a measure of the historical return difference between a fund and its benchmark. Tracking error looks backward Using returns from a portfolio and an index:

$$TrackingError = \sqrt{\frac{\sum_{i=1}^{N}(R_{fund} - R_{index})^2}{N-1}}$$

where R_{fund} is the return of the fund, R_{index} is the return of the comparison index, and N is the number of return period observations.

The difference between the two measures is that tracking error looks backward at returns while Active Share is calculated from the differences in current holdings between a fund and its benchmark. Active Share is represented as a more accurate way of measuring manager activity rather than simply looking at ex post tracking risk. This is because a high stock concentration manager who diversifies across industries will look less active than a diversified manager who rotates holdings among sectors.

The key claim of Active Share is that it predicts relative performance and is robust against factors such as turnover, expense ratio, and the number of stocks in a portfolio. In addition, the Carhart alpha model (FF4) does not explain any significant relationship between Active Share and Performance.

Petajisto updated their work (Petajisto 2013) to incorporate the history of the 2008 financial crisis period, extend the sample size of funds examined, and point out more explanations for why this measure predicts performance. In addition, a new set of analysis is introduced using market volatility as a potential explanatory variable for performance.

While the idea of Active Share seems intuitive, the concept came under heavy attack from (Frazzini, Friedman, and Pomorski 2016) who find that the empirical evidence for the measure is not very robust. The difference in out-performance between high and low Active Share funds is driven by the strong correlation between Active Share and the benchmark type and does not predict actual returns. Within individual benchmarks, Active Share is as likely to correlate positively with performance as it is to correlate negatively. Therefore, the authors conclude that Active Share is not useful for selecting managers.

The debate raged on over this measure of manager activity and the forecasting of skill based returns with a rebuttal to the criticism and the publication of (Cremers 2017). In this further demonstration of the power of Active Share, the author lays out three pillars of active management as the triad of knowledge, judgment, and effective application. The authors do concede that even among high Active Share managers, only those with long holding periods tend to create superior performance.

While Active Share does seem to measure how different a portfolio is relative to its benchmark, it does not answer the question posed by this thesis. How does the capital allocation method interact with security selection in pursuit of superior portfolio efficiency? Active Share can not answer this question because it entangles the two stages of portfolio formation.

1.5 Methods of Performance Measurement

The application of active management creates inevitable performance differences. The aim of active management is to 'out-perform' the comparison portfolio. Since managers know that their performance is being measured relative to a benchmark, there is the possibility that these measures can be manipulated as was shown by (Goetzmann et al. 2007). The author finds that while moral hazard exists, to measure performance fairly the metric should be manipulation proof and not reward information free trading.

1.5.1 Choosing the Benchmark

For a measure of relative performance to be meaningful, the choice of comparison portfolio should also be appropriate. It makes no sense to compare a bond portfolio to a stock portfolio and claim victory if performance is better. But what are the right ways to compare? Several methods have been introduced in the literature.

Early in the literature, (Treynor 1965) showed a simple early version of how to rate the management of investment funds, noting early on that the choice of comparison is extremely important. This importance is echoed in (Grinblatt and Titman 1989) who caution about the possibility of overestimating risk due to market timing abilities and the failure of investors to earn positive risk adjusted returns due to increasing risk aversion. The authors of (Grinblatt and Titman 1994) show in a mathematically rigorous manner that the choice of benchmark is incredibly important to measuring success and that when a mean-variance inefficient benchmark is used, erroneous inferences arise.

1.5.1.1 Style Benchmarks As (Brown and Goetzmann 1997) explain, styles are buckets that are used to characterize managers. Benchmarking performance is seemingly more relevant when managers of the same style are compared against each other, but style consistency is required. There is complexity in this process however as the Investment Company Institute (ICI) classifications are too broad to carry much meaning. Allowing managers to select their own styles can lead to moral hazard and incorrect choices and allows managers to change their style benchmark election through time to game the measure of relative performance.

For a style benchmark to be meaningful, it must be assigned objectively and empirically. The authors provide an algorithm to aid in classifying funds by style using likelihood ratio tests. Tests of their methods show significant improvement over manager selected styles.

An extension of the style model builds a reference portfolio from the characteristics of the portfolio being examined. In (Daniel et al. 1997) the authors construct benchmarks that match the fundamental characteristics of the stocks in the subject portfolio using market capitalization, book to market ratio, and prior year returns. Additional 'characteristic timing' and 'characteristic selectivity' measures are developed to detect these skills in the subject portfolios.

1.5.1.2 Characteristic Benchmarks While the factor approach to benchmarking manager alpha used regression of historical returns, the work of (Daniel et al. 1997) introduced a slightly different method. The characteristic approach uses the fundamental characteristics of the current holdings to map to a set of benchmark holdings. The goal is the same as with factor models, which is to isolate the true value added by the manager.

1.5.1.3 Factor Models With the addition of momentum, (Carhart 1997) moved the work forward in explaining the persistence in mutual fund performance. The factor models helped to push forward the explanation of returns and create a new benchmark for manager performance. Rather than using just market returns to explain fund performance, the factor models extended the analysis to a multivariate regression. This minimized the unexplained error term α, which is assumed to be the value added by the manager in a particular period. Factor models have shown to be useful in measuring the performance alpha of a manager.

1.5.2 Choosing the Performance Metric

Markowitz's second tenet solidly rejects return as the only metric for judging the success of a portfolio, arguing that without perfect foresight, diversification is necessary to achieve an acceptable outcome.

> *"The hypothesis (or maxim) that the investor does (or should) maximize discounted return must be rejected." "Diversification is both observed and sensible; a rule of behavior which does not imply the superiority of diversification must be rejected both as a hypothesis and as a maxim." (Markowitz 1952)*

A simple return comparison is not appropriate as the portfolio construction process was a choice based on risk and return, so it is only logical to also have a two dimensional measure for inspecting the success of the process. But which metric of the multiple available should be used?

The work of (Chen and Knez 1996) states clearly that any admissible portfolio performance measure should satisfy four minimal conditions: assign zero performance to the reference portfolio, be linear, continuous, and nontrivial. These conditions can only exist in a market free from arbitrage.

Some examples of two dimensional performance metrics are:

$$SharpeRatio = \frac{r_i - return_f}{\sigma_i}$$

$$TreynorRatio = \frac{r_i - return_f}{\beta_i}$$

where r_i is the return of the investment, r_f is the return on the risk free rate, σ_i is the standard deviation of the returns of the investment, and β_i is the slope coefficient of the linear regression between the returns of the investment and the returns of the market.

Another example is the Calmar ratio which uses the rolling 36 month annualized returns in the numerator and the maximum drawdown in the denominator. Whereas the first two metrics above measure performance relative to an external benchmark, the Calmar measures performance relative to its own path.

Many other metrics exist which incorporate other components such as draw-down, skewness, partial deviations, etc. It is the responsibility of the investor to determine which set of metrics most closely resemble their risk aversion parameters and then calculate and compare these metrics.

1.6 Performance Attribution

Finally, there exists a body of literature related to performance attribution. Robust explanations of performance have been documented by (Fama 1972), (Bollen and Busse 2001), (Brinson, Hood, and Beebower 1995) and others which explained excess performance variability by decomposing it into elements such as policy, timing, and security selection. Three additional works of importance in this area are (Ibbotson 2010), (Xiong et al. 2010), and (Ibbotson and Kaplan 2000). While their work is helpful, in each case, these works focus on portfolios containing multiple asset classes and are an ex-post explanation of excess return while my research is focused on documenting the interdependence of security selection and capital allocation method, and examining the drivers of optimal Sharpe ratio portfolios.

1.7 Can Managers Beat the Market?

For all the funds that exist and all the various ways that managers can apply components of active decision making, the big question remains: can managers beat the markets? As we reviewed in the previous section, the answer depends on the reference portfolio and the statistic used for comparison.

In one of the earliest studies of manager performance, (Sharpe 1966) noted that each investor must exhibit their own preferences about the trade-off of risk and return and that a manager cannot directly observe the preferences of investors. All that the manager can do is establish a fund and invite investors with similar preferences to join.

Following Sharpe, (Jensen 1968) also examines the performance of a group of early mutual fund managers and finds no evidence of a managers ability to pick stocks and that performance holds very little evidence that any individual fund was able to do significantly better than that which we expected from mere random chance. His findings hold even in the case of gross returns.

In the following period, (Malkiel 1995) updates the analysis with a larger database at a later date range and finds similar conclusions but adds that strategies built on buying the 'hot hand' also do not work well.

Whether the data support the 'Yes' or 'No' answer, the authors of (Baks, Metrick, and Wachter 2001) suggest that a returns based analysis may be too simplistic and that a Bayesian approach which incorporates more data points such as cash flows into the fund, age of the manager, education, SAT scores, etc. With additional data points, even very skeptical prior beliefs can lead to rational and economically significant allocations to active management.

1.7.1 Maybe 'No'

The majority of the literature seems to document the inability of active managers to outperform. The work of (Carhart 1997) is commonly cited. Carhart created a unique data set which corrected for survivorship bias in the returns databases. His finding is that common factors in stock returns and investment expenses almost completely explain persistence in equity mutual funds' mean and risk adjusted returns.

Carhart developed an extension of the Fama and French factor model (FF3) by incorporating momentum as an additional factor (FF4). With his FF4 model he shows that persistence in fund performance alpha can not be seen in the data and that the turnover of funds in the top quartile of return performance is roughly 80% each year. The only significant persistence not explained by these common factors was with the worst performing funds. Carhart's conclusion was that "the results do not support the existence of skilled or informed mutual fund portfolio managers."

With a decomposition of returns across the categories of selectivity, timing, and style, (Wermers 2000) showed that managers create approximately 1.3% per year from stock selection, lose 0.7% due to non-stock holdings, and lose another 1.6% due to transaction costs and expenses. The net loss is approximated at 1% per year for the industry while at the same time supporting the claim that managers can add value through stock selection.

In three important papers, researchers looked to differentiate skill from luck. The authors of (Cuthbertson, Nitzsche, and O'Sullivan 2008), (Fama and French 2010) and

(Barras, Scaillet, and Wermers 2010) identify managers in the categories of skilled, unskilled, and lucky. In general, the performance of the group of all managers have a performance that is very similar to the total market. They acknowledge that since there are other investors who are not in the funds, that overall it is not appropriate to conclude that fund investing alpha is a zero-sum game. However, less than 1% of managers show persistent performance that is distinguishable from luck.

In (Berk and Binsbergen 2015), the authors look at performance from a dollar perspective so that larger funds have more weight in the analysis of the question. In their analysis, the average fund generates $3.2 million per year of skill-based alpha which can persist for as long as ten years. Measuring skill via dollars extracted is a way to neutralize the effect of declining alpha on increasing assets under management. The authors note that an important economic fundamental is that agents earn rents only if they have a competitive advantage. However, as money managers are some of the most highly paid professionals, this tenet seems to be violated as net of costs, this group of managers shows no skill and any alpha created is extracted by high costs.

In a very different view on under-performance, the authors of (Glode 2011) cite the logic of investing in under-performing funds. Asset pricing theory states that the return premium is not a function of variance but of covariance. So investing in funds that underperform in general are economically consistent if their covariance in negative states of the world is lower. Mutual fund investing and negative expected fund performance can simultaneously arise in a setting with skilled fund managers facing rational investors.

1.7.2 Maybe 'Yes'

A minority of the literature documents some managers who outperform. The Active Share authors in (Cremers and Petajisto 2009) and (Petajisto 2013) show persistent out-performance as the Active Share Measure increases, although this was thoroughly challenged in the practitioner literature.

Managers who can out-perform apply a set of time varying skills as shown in (Kacperczyk, Nieuwerburgh, and Veldkamp 2014). Stock picking is evident in booms and market timing is evident in recessions. Even more important is that the same managers that show proof of stock selection skills also display market timing abilities. Unconditional returns of these top quartile skilled managers is 0.50% to 0.80% per year better than other funds. To emphasize this ability, the authors suggest a new measure that weights stock picking skills during recessions and weights market timing skills during recessions which are less common.

Using a bootstrap analysis of the joint distribution of performance measures (alphas) across all funds to extract luck from skill, the authors of (Kosowski et al. 2006) find that a sizable minority of managers really do pick stocks well enough to more than cover their costs. Bootstrap analysis indicates that the superior alphas of these managers do persist, but are highly concentrated in growth oriented funds.

1.7.3 Performance Inconclusive

At a high level, it is impossible to conclude from the literature that professional money managers have been able to consistently apply their expertise to regularly add return to the investment process. We do see various ways in which pockets of alpha are created by some managers, but between the effects of high economic rents, and the inconsistency of the performance, it seems nearly impossible for investors to have chosen how or where to invest their money in actively managed strategies.

Could it be that professional managers are applying their resources less effectively than possible? This thesis will quantify the relative effects of security selection and capital allocation methods which may create opportunity for managers to apply their resources more effectively.

1.8 Mechanics

Research leading to a clear conclusion requires the application of a careful methodology to clean data and analyzed using the most applicable statistics in a robust manner. This is not a simple undertaking. As we see in the progression of techniques in the time series of the literature, much has occurred in the last 50 years to allow for this type of analysis to be completed. Today we are in a new age where massive computational power is available to researchers for a tiny cost. This is allowing for old questions to be analyzed in new ways, and this thesis will take advantage of these previous mechanical advances and this new computational leap forward.

1.8.1 Data Availability & Quality

By mandate of the US Securities and Exchange Commission (SEC), the holdings and returns for US Equity Mutual Funds has been the most widely available data source for studying the effects of active management. The consistency of the data and the combination of data elements needed to perform empirical testing was not always present in the form that now exists.

A breakthrough of collecting, cleaning, and correcting bias issues in the data was from the work of (Carhart 1997) who brought together data from January 1962 to December 1993 using Micropal/Investment Company Data, Inc. (ICDI), FundScope Magazine, United Babson Reports, Wiesenberger Investment Companies, and the Wall Street Journal. Survivorship bias is a major issue in the raw data as by 1993 about one third of total funds in the database had ceased operations.

An analysis of the quality of the data was completed more recently by (Elton, Gruber, and Blake 2001) who found that the CRSP return data still retains upward biases and that data from manager mergers are inaccurately recorded about half of the time. Differences in returns between CRSP and Morningstar are a problem for older data and smaller funds. While work to remove survivorship bias has been successful, errors of omission are also significant and have the same effect. In the end, all data sets have errors but it is important to recognize and correct for known systematic errors.

1.8.2 Methods

A wide array of analytical methods is seen across the literature. A sampling of the categories and approaches is represented here.

1.8.2.1 Portfolio Formation In (Chow et al. 2011) and similar approaches, portfolios are formed independently using a disclosed set of rules. An example of this would be selecting the 500 largest capitalization stocks and forming a portfolio for analysis. This approach differs from simply using the standard S&P 500 which may have slightly different constituents and weightings due to corporate actions, etc. but lends itself well to reproducible research.

1.8.2.2 Rebalancing Some research uses single periods of rebalancing, such as annual, for their analysis. Other approaches include (Maillard, Roncalli, and Teiletche 2008) who build their back-tests on a rolling sample approach and rebalanced each month.

1.8.2.3 Data Sets & Metrics The very rigorous work of (DeMiguel, Garlappi, and Uppal 2009) go beyond a range of formation and rebalancing methods by extending their work with multiple metrics and applications to an array of domestic and international equities.

1.8.2.4 Parameter Estimation Where sample estimates are used as inputs for optimization, the better analytical research pieces vary the parameters for calculation. For example, in (Ledoit and Wolf 2004), the authors use a variety of input parameters for testing their approach.

A full review of analytical methods is beyond the scope of this literature review. However, it is important to have a sense of the array of techniques and the tendency of the most rigorous research to employ an array of these techniques.

1.9 Methods of Robustness

Research conclusions are based on observations of the data and the statistics calculated. The best quantitative methods are applied to glean insights which are reasonably correct. This becomes difficult when the calculations can be accomplished in multiple ways or the output of the calculations are too close to know for certain if a difference is probable. Fortunately in the past 50 years, there have been researchers who focused on improving the methods of robustness, and today the best research work incorporates these new tools.

1.9.1 Construction Tools

Many approaches to portfolio construction rely on estimated parameters such as volatility and correlation measures. When these parameters are used within an optimization, errors in the estimation are magnified in the optimization output. While there are several methods for smoothing these errors, the standard approach appears to be (Ledoit and Wolf 2004) which introduces a shrinkage process for minimizing estimation errors. This approach is now widely followed throughout the practice and is referenced often in the literature.

Another interesting view into this issues is (DeMiguel and Nogales 2009) which was published in the Journal of Operations Research. The authors cite the important work of (Merton 1980) who had shown that the error introduced when estimating mean returns is much smaller than the error from estimating the covariance matrix. To combat this issue, (DeMiguel, Garlappi, and Uppal 2009) introduce robust estimators. The unique contributions of this work was to show how to compute the portfolio policy that minimizes a robust estimator of risk by solving for a single nonlinear program with a focus on minimum variance. Previous research had only been able to show this solution in a two step approach to the calculation.

1.9.2 Performance Verification

Performance measurement statistics (e.g. Sharpe, Treynor, or Calmar ratios) are typically point estimates made from a time series of returns. Comparing two point estimates however is not conclusive. The statistic being analyzed is a sample estimate of the true statistic and therefore it has a range of confidence. It may be that two investment managers with different performance measures are statistically indistinguishable due to the confidence bands needed around the sample statistic.

The work to understand this effect was originally published by (Jobson and Korkie 1981) where the authors looked at the Sharpe and Treynor ratios. They found that for comparisons between just two managers, the Sharpe $z - statistic$ is well behaved but if multiple managers are compared, the χ^2 statistic is more appropriate. Interestingly, they also found that the Treynor measure is not well behaved and is therefore unsatisfactory for comparisons. While the work was thorough, corrections were made by (Memmel 2003).

In a further work, the authors of (Ledoit and Wolf 2008) pointed out that the previous tests are not valid if the distributions have heavier tails than normal or are of a time-series nature. In this frequent case, robust estimators can be constructed by forming a studentized time series bootstrap confidence interval for the difference in Sharpe ratios. If zero is not contained within this interval, than it can be declared that the two ratios are different with a high degree of confidence. The method introduced by the authors is used extensively in my research.

High quality research relies on high quality methods and calculations. Use of a robust performance verification method is critical and will be implemented within this thesis.

1.10 Contribution to the Literature

Constructing an optimal portfolio for tomorrow and beyond is an impossible task without perfect foresight or luck. Yet professional managers are persistently employed to apply their expertise despite widespread evidence of under-performing even simple representations of broad markets.

A survey of the existing literature reveals a great set of work on portfolio construction, yet only a small set of literature exists which compares the relative effectiveness of security selection and capital allocation within the portfolio formation process. Of the few pieces, (Maillard, Roncalli, and Teiletche 2008) does examine whether superior returns from systematic capital allocation can still exist when controlling for factors such as the small cap bias. The authors do some introductory work in the relative importance of these two decision steps. Beyond this, very little is found to answer my particular questions.

In a series of four papers, I examine the relationship between security selection and capital allocation in new ways. While the nature of the problem does not lend itself to deriving a closed form answer, looking at the issue from four different perspectives allows for a deeper understanding and hopefully leads to more efficient allocation of professional resources. By addressing this gap in the literature, it is hoped that my work serves as a base for future research while creating improved portfolio efficiency for investors in the future.

Chapter 2

The Hierarchy of Sharpe Ratios

2 The Hierarchy of Sharpe Ratios: Observed vs. Significant

2.1 Introduction

The purpose of this thesis is to examine the interconnected nature of security selection and capital allocation, describe the relationship, and then to explore the drivers of optimal Sharpe ratio portfolios. This chapter takes the first step toward untangling this connection between selection and allocation. I start by holding security selection as constant and robustly examining the benefits of systematic capital allocation. In works cited in my literature review, such as (Clare, Motson, and Thomas 2013; Arnott et al. 2010; Martellini 2008) and others, there are purported benefits that systematic capital allocation methods have over the market capitalization method. In subsequent chapters, security selection is introduced as a variable which helps to further untangle the relationship between the two stages of portfolio formation.

2.2 The Opportunity Sets: Security Selection vs. Capital Allocation

There are two stages of portfolio formation according to Markowitz. Each of these potential opportunities to improve a portfolio are summarized in Figure 1 and Figure 2.

The opportunity set for security selection within the largest 1,000 stocks each year appears enticing as shown in Figure 1. On average, the yearly range of performance between the top and bottom quartile of stocks is nearly 40%. And while this varies over time, the normal range is wide enough to consider each year a potential 'stock picker's market'.

Selecting the best performing stocks and avoiding the worst performing stocks from this wide range would seem like a rich source for harvesting alpha. In addition, we see that

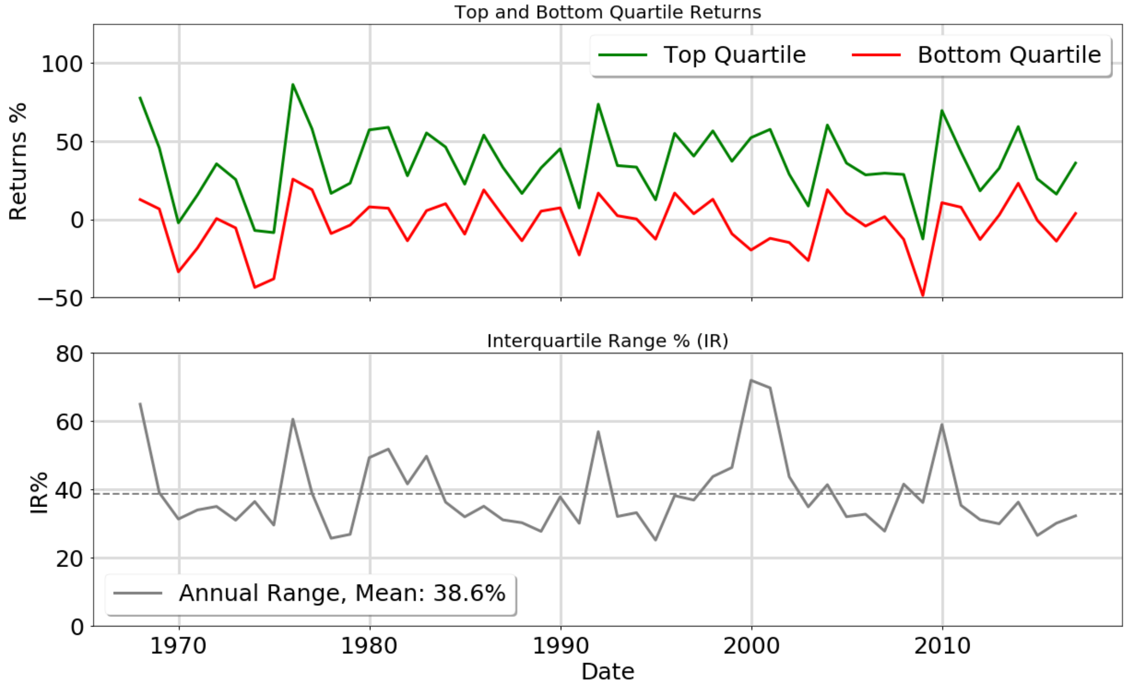

Figure 1: Annual Returns - Top 1,000 Largest Stocks
The top panel displays the 25th and 75th percentile returns for the 1,000 largest stocks by market capitalization in the data set each year. The bottom panel displays the inter-quartile range (top quartile - bottom quartile) of returns each year.

the distribution of equity returns in any particular year is highly and positively skewed as stocks can only drop by 100% while they can rise an unlimited amount. Picking winners would seem to more easily offset picking losing stocks.

The opportunity set for portfolio diversification as shown in Figure 2 is more difficult to discern as its benefit is non-linear. The data set shows that mean monthly return correlation between stocks in the 1,000 largest capitalization opportunity set have remained low and stable around 0.27 which provides a rich source for creating diversification benefits to a portfolio. A few studies have also focused on the potential for stability in the hierarchy of the correlation structure as an additional diversification benefit that can be harvested.

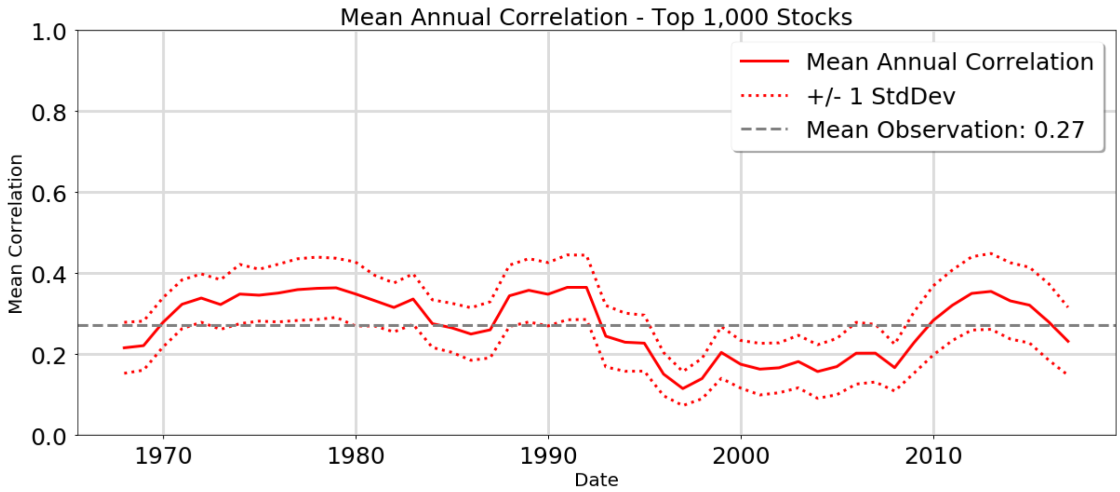

Figure 2: Mean Annual Correlation - Top 1,000 Largest Stocks
The figure above shows the mean correlation each year between the top 1,000 stocks as well as the mean observation of all years with a 1 standard deviation band. The standard deviation band is calculated at each year in the time series and is intended to show the tight range around the low mean correlation.

From a naïve perspective, security selection may appear to be an easier source of potential alpha as higher returns from the stocks selected will create higher returns for the portfolio overall. However, good stock picks may be offset by bad stock picks and the net effect of security selection alpha is easily neutralized. Returns each year and across individual equities are inconsistent, but low and stable correlations in the cross section of returns are a reliable method to create portfolios with lower volatility which manifests in higher risk-adjusted returns (e.g. Sharpe ratio) across time.

However, while there is no known source of employment data disclosing the count of equity analysts vs. portfolio managers, active equity managers appear to spend a disproportionate share of effort on picking stocks. Typically with an active equity manager there are multiple analysts for each portfolio manager.

The skills required to operate in each of these opportunity sets are disparate. The skills of fundamental security analysis are quite different from the mathematical skills required to understand the complexities and embedded assumptions in various capital allocation methods.

A possible explanation for the dominant focus on security selection may be that the familiarity of the fundamental analysis process is carried over during the promotion from analyst to portfolio manager. Another possible explanation is that the mathematics of portfolio optimization methods are outside the familiar business school training of many investment professionals.

When given responsibility for both security selection and capital allocation, portfolio managers tend to rely on their 'best ideas' that have emanated from the security selection stage rather than harvesting from the rich source of diversification benefit which is embedded in the capital allocation methods. Another explanation may be that 'stock stories' are more easily understood. Explaining performance due to Apple vs. Lehman Brothers is more comprehensible than an explanation related to an unanticipated shift in the hierarchy of correlations.

Moving beyond a capitalization weighted allocation method to an alternative method which emphasizes the benefit of diversification would seem to be a value added approach, but as we will show, selecting the optimal allocation method is not as clear as picking a method with a higher ex-post mean Sharpe ratio. By applying a robust bootstrap method to examine the differences in Sharpe ratios, the hierarchy of performance improvement is less clear.

2.3 Selection vs. Allocation

But first, it is important to clarify the meaning of the terms 'security selection' and 'capital allocation'.

These two terms refer to the two stages of portfolio formation as described by (Markowitz

1952). Security selection is the process of beginning with a full set of available investment choices, and through some process narrowing this set down to a smaller subset for potential inclusion in the portfolio. The reduction from the full set may be accomplished in a variety or combination of methods. For example, it could be narrowed by characteristics of the stocks such as capitalization size, country, industry or sector, or factor exposure such as growth vs. value. The narrowing of the set may also occur by discretionary method such as fundamental valuation techniques applied to company financial statements or interviews with management. It could also be accomplished in a systematic method such as screening for trend/momentum or other quantitative characteristic. And combinations of these methods are also available. An example of a combination is a 'quantamental' manager who uses quantitative methods applied to fundamental characteristics to systematically select stocks. The potential methods of security selection are endless and are a reason why I later introduce a method to select securities as a limit condition that bounds all of the potential methods contemplated here.

The second term is 'capital allocation' which in certain parts of the world may be referred to as a 'weighting scheme'. The meaning of the term is the method by which available capital is spread across the securities that were selected in the first stage of portfolio formation. Again, there are numerous ways to accomplish this step. Purely discretionary methods are sometimes called 'conviction weighting' and are carried over from the sense of certainty that arose during the security selection stage. Other methods are systematic and include many of the methods studied in this thesis, such as equal weight, minimum variance, etc.

2.4 Capital Allocation Methods

2.4.1 Methods Available But Not Chosen

Two comprehensive reviews of the variety of allocation methods are (Chow et al. 2011) and (Clare, Motson, and Thomas 2013). These studies document a wider variety of allocation methods than are examined in this thesis. The additional methods listed here were not selected for study simply because in the time that in the 10 years plus that they have been documented, none of them have gained any substantial existence in commercially available investment funds. While there are many potential reasons for their lack of popularity, the practical nature of my questions led me to constrain the panel of methods examined to those that have gained traction in practice.

The additional methods observed in the literature, but not selected for my panel, are:

- Risk Cluster Equal Weight - the basic equal weight method is subject to the stock universe selected and may end up with an over allocation to a set of stocks that is tightly clustered in risk space. This method attempts to overcome this issue by clustering stocks based on their correlations and then equal weighting the clusters.
- QS Investors' Diversification Based - this method moves a step further than the previous method by first bucketing stocks into countries and sectors and then equal weighting the buckets. While the simple equal weighted portfolio can show significant changes in the risk and return characteristics of the resulting portfolio, an advantage of this method is that it more closely matches the risk and return characteristics of the starting capitalization weighted reference portfolio.
- Diversity Weighting - Alternative allocation methods can produce significant tracking error against a cap weighted reference portfolio. The diversity weighting

method solves for this through an interpolation method between the two allocation methods with a focus on tracking error.

- Fundamental Weighting - there is a nearly limitless set of methods that can be derived under this banner. But in general, fundamental weighting takes characteristics of the underlying companies (such as sales, book value, etc.) and uses them to systematically weight the securities. This approach can create a wide variety of allocation methods depending on the set of fundamental characteristics used in the blending approach.

- Mean Variance Optimization - this method, set in place by (Markowitz 1952), was the foundation to portfolio formation. Before this method, there did not exist much in the way of systematic allocation methods. The approach uses estimates of return, volatility, and correlation to form an efficient frontier from which an investor could choose a most efficient portfolio. Researchers have pointed out several inherent issues with the methodology which likely prevent it from being widely used in practice today. From the array of concerns raised, there are two main issues that rise to the top. First is that this optimization method maximizes the errors of the estimates. Very small differences in any of the parameters can create vastly different resulting portfolios. The second main issue is that portfolios found along the frontier can be highly concentrated, with very few positions. These are likely the reasons why this method is not widely observed in practice today. Even for those who employ the method as a starting point, constraints on concentration limits are typically used and these are often binding and cause the resulting frontier to be substantially below the unconstrained frontier.

- Maximum Sharpe Ratio version 1 (MSRv1) - this method attempts to improve upon the minimum variance portfolio which does not incorporate expected returns. For minimum variance to be the optimal portfolio, all stocks must have the same expected returns. As this is not likely, the MSRv1 incorporate an expected return based on a static ratio of return per unit of volatility at the stock level. With this additional assumption for the return input, the MSRv1 optimizes for Sharpe ratio.

- Maximum Sharpe Ratio version 2 (MSRv2) - this method follows the same reasoning as MSRv1 but rather than assume stock level returns are a static relationship to the full distribution of volatility, the method ties expected return for each stock to the level of downside volatility. Again, with a new assumption for expected returns, this additional input is used to optimize for the maximum Sharpe ratio.

- Risk Efficient - this method is closely related to MSRv2. Stocks are clustered into deciles based on downside volatility, and then the median downside volatility of each decile is used to proxy for the expected returns in that decile. Once all deciles are measured for median downside volatility and the stocks in the deciles have proxied expected returns, then the portfolio is optimized to find the portfolio with the highest expected return relative to the portfolio standard deviation. This may still create a very concentrated portfolio of stocks and so the method imposes further restrictions with regard to individual stock constituent weights.

- Random Weight - this method of weighting stocks is often seen in the literature as a simple comparison against systematic methods. Random weighting can be achieved in an unlimited number of ways and so does not lead to a method which can be simply documented and applied. This makes the approach helpful for a research based approach, but is not practical in application. Clearly there are no commercially available products based on this method.

- Systematic/Random - an advance on the purely random methods, this set of methods begins with an assignment of values to the stocks which have no true mathematical or sensible link. For example, Nick Motson describes a Scrabble weighting scheme which assigns Scrabble scores to the letters in the tickers of the stocks.

Then the scrabble score is summed for each stock and weights are assigned as the proportion of the stock score to the total score of all available stocks. Other systematic/random methods could certainly be created. And while these methods don't contain an accessible set of fundamental logic, they are useful for checking whether the benefits observed in any of the other methods are significantly different from a random approach. This gives these methods a value in helping to verify what might be observed in more fundamentally logical approaches. Methods in this category are never observed in practice and so have been excluded from examination in my panel of allocation methods.

2.4.2 Capital Allocation Methods Chosen For Examination

While the literature describes a wide variety of systematic allocation methods, the panel chosen for this study is a subset of methods that are commercially more available today. The allocation methods examined in this thesis are: market capitalization (MC), equal weight (EW), inverse volatility (IV), equal risk (ER), minimum variance (MV), and maximum diversification (MD). In the existing literature, the MV and MD methods are optimized using both 1% and 5% maximum position size limits. I have followed this approach and the two variants of the approaches are denoted as MV1, MV5, MD1, and MD5.

To conduct this analysis, it is necessary to rebuild much of the data and methods used in previous literature. The methods and data used for this work are described fully in Appendix 1. Each of these portfolios are formed annually using the 1,000 largest capitalization US equities. A full explanation of the data and methods for constructing these portfolios are described in the appendix of this thesis as well as a set of summary statistical characteristics which is shown in Table 17 in Appendix 1.

While there are documented improvements in Sharpe ratio from alternative allocation methods, 'not everything that glitters is gold'. As a starting point, the purpose of this chapter is to more clearly define the hierarchy of Sharpe ratios across a panel of allocation methods and to examine their robustness.

Each of the selected allocation methods was tested across the full 50 year history (1968 - 2017), with the set of available stocks each year being the 1,000 largest market capitalization companies in the database. A summary of the results and characteristics is displayed here in Table 1. Each of the performance and risk statistics summarized in the table is explained in detail in Appendix 1 of this thesis.

2.4.3 Description of Capital Allocation Methods

It has been well documented in previous studies such as (Chow et al. 2011) and (Clare, Motson, and Thomas 2013) that following any one of a range of alternative capital allocation methods would have led to superior returns and/or return efficiency. In these previous studies, the authors analyzed a wide range of methods and parameters. My study begins by selecting a panel of established and practical allocation methods and focuses on documenting the drivers of Sharpe ratios. To maintain focus on the interaction effect, a smaller set of allocation methods have been chosen which are more prevalent commercially and have wider practical appeal.

For purposes of clarity, the term used in this paper is 'capital allocation method' which has the same implied meaning as 'weighting scheme' which is a more common term in some parts of the world. In either form, this step in the portfolio formation process is concerned with spreading the available money across the securities which have been selected for the portfolio.

2.4.3.1 Market Capitalization Market capitalization weighting is the most commonly available form of allocation method for investing passively and this is seen in the major market indices such as the S&P 500. While this allocation method is shown historically to have returns which are inferior to other allocation methods, it is highly efficient in a practical manner. As market capitalization of each security changes due to price movements each day, the weights change in unison requiring no trading turnover until the index is reformed.

The CRSP database provides share price and the number of shares outstanding each month. These two elements are multiplied together to create the market capitalization for each stock each period.

For consistency across the panel of allocation methods, I form a market capitalization weighted portfolio from the 1,000 largest securities available each year rather than use a market provided index which would have slightly different securities than would be used in my other allocation methods. This is an approach that is observed in the literature. The weighting for each stock (w_i) is found by dividing the individual stock's market capitalization by the sum of all stock's market capitalization on the date of index formation.

$$w_i = \frac{MktCap_i}{\sum_{i=1}^{N} MktCap_i} \tag{1}$$

The weights for the index are calculated using data from the last day of each year and the index is reformed and held without rebalance for the following year until the process is repeated. While this is not necessarily the exact selection and rebalance method that commercially available indices such as the Russell 1,000 implement, it is a method that can be applied consistently across a panel of allocation methods for this study and is consistent with methods employed in existing literature.

2.4.3.2 Equal Weight

In a well documented study, (DeMiguel, Garlappi, and Uppal 2009) show the benefits and simplicity of an equal weight allocation methodology. By following an equally weighted method, there are no parameters to estimate, and the weight calculation is trivial as it is simply the inverse of the number of stocks selected for inclusion in the index.

$$w_i = \frac{1}{N} \tag{2}$$

In addition to describing the simplicity and robustness of the equal weighted method, the authors show the period of historical data that would be required in order to accurately estimate parameters for other methods in order to outperform the simple equal weighted approach. A sample mean variance optimization based portfolio with 50 assets would require 6,000 monthly observations in order to accurately estimate the covariance matrix and outperform the simple equal weighted method.

One note of caution is provided by (Chow et al. 2011). The authors state that "A notable feature of equal weighting is that the resulting portfolio risk-return characteristics are highly sensitive to the number of included stocks. Although the S&P 500 Index and the Russell 1,000 Index have nearly identical risk-return characteristics over time, the equal-weighted S&P 500 portfolios and the equal-weighted Russell 1,000 portfolios have dramatically different risk-return characteristics."

2.4.3.3 Inverse Volatility

The inverse volatility allocation method attempts to distribute risk more equally across the set of available stocks while estimating a single parameter, volatility. As a result, one measure of diversification (effective number of stocks) is increased dramatically over the market capitalization method.

The method also attempts to capture a documented low-volatility anomaly whereby stocks with lower volatility have been shown to have higher risk adjusted returns. Building on their previous work from 1972, (Baker and Haugen 2012) show in a simple and transparent way that this anomaly has presented itself across long periods of time (US 1926 - 1990) and across the world (33 global stock markets during 1990 - 2011).

To create the weights for each stock in my study, the standard deviation of each stock (σ_i) is estimated using the previous 5 years of monthly returns. By using the inverse of this measure ($\frac{1}{\sigma_i}$), this allocation method shifts weight from stocks with higher volatility to stocks with lower volatility. The weight for each stock then is given by:

$$w_i = \frac{\frac{1}{\sigma_i}}{\sum_{i=1}^{N} \left(\frac{1}{\sigma_i}\right)} \tag{3}$$

The inverse volatility method combines increased diversification with harvesting the low-risk anomaly. While the method uses a single estimated parameter (standard deviation), the assumption under which this method would produce the optimal risk adjusted portfolio is that correlations are uniform across all securities.

2.4.3.4 Equal Risk Contribution When the uniform correlations assumptions required for optimality of the inverse volatility method are not realized, then groups of highly correlated stocks receive larger allocations of risk. By including correlations between stocks into the allocation process, the equal risk contribution method accounts for the non-uniform correlation structure.

Estimating a second parameter introduces additional risks to the calculation. Calculating the weights which provide an equal contribution to risk from each stock does not have a closed form solution and must be solved using optimization. Optimization methods tend to amplify estimation errors, and this method now encounters estimation problems, especially as the covariance matrix is highly dimensional with up to 1,000 stocks.

To minimize the effect of these estimation errors, the analysis employs the shrinkage method described in (Ledoit and Wolf 2004) which condenses the outlier observations toward the central values in a systematic way and reduces the amplification of errors produced by the optimizer. The calculations are performed in Python using sklearn.covariance.LedoitWolf
Use of this shrinkage method is extended for all allocations in my research that require a covariance matrix.

All of the allocation methods examined in this study require long only positions ($w_i \geq 0$) with a constraint to be fully invested $\left(\sum_{i=1}^{N} w_i = 1 \right)$. To solve for the weights which produce an equal contribution to risk of each stock we follow the method described by (Maillard, Roncalli, and Teiletche 2008) and first define the minimization function.

Let w be the vector of weights and let Σ be the covariance matrix. Then the volatility of the portfolio is defined as:

$$\sigma(w) = \sqrt{w' \Sigma w} \tag{4}$$

The risk contribution of each asset within the portfolio is given as:

$$\sigma_i(w) = w_i \times \partial_{w_i} \sigma(w) \tag{5}$$

[2] https://scikit-learn.org/stable/modules/generated/sklearn.covariance.LedoitWolf.html

From this is constructed the vector of marginal contributions ($\partial_{w_i}\sigma(w)$) as:

$$c(w) = \frac{\Sigma w}{\sqrt{w'\Sigma w}} \tag{6}$$

The solution for the weights is given by the optimization:

$$\underset{w}{arg\,min} \sum_{i=1}^{N} \left[\frac{\sqrt{w'\Sigma w}}{N} - w_i \cdot c(w)_i \right]^2 \tag{7}$$

The minimization routine for this analysis was constructed with the Python 'SciPy' package using the sequential least squares programming (SLSQP) method.[3]

2.4.3.5 Minimum Variance In a traditional mean-variance optimization approach, three parameter estimates are required: returns, volatilities, and correlations. Returns are the most difficult to forecast and errors in the estimated returns create large errors in an optimization output even when using a covariance shrinkage model. If it can be assumed that all stock returns are the same, then the optimal solution can be found by estimating just the volatility and correlation (as in the equal risk method) and solving for the weights that minimize the risk of the portfolio.

Depending on the parameter estimates, this allocation method may create a very concentrated set of weights. To prevent this possibility of concentration, my analysis follows the method of previous academic studies in limiting the maximum position size to 1%. For comparison purposes relative to the maximum diversification approach, I also include maximum position sizes at 5%. As above, this analysis requires long only positions ($w_i \geq 0$) and constrained to be fully invested ($\sum_{i=1}^{N} w_i = 1$).

To solve for the weights which produce a minimum variance portfolio, the minimization

[3] https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html

function is defined so that w is the vector of weights and Σ is the covariance matrix. The optimal weights are found by:

$$\min_{w}(w'\Sigma w) \text{ subject to } \begin{cases} \sum\limits_{i=1}^{N} w_i = 1 \\ 0 \leq w_i \leq max\% \end{cases} \tag{8}$$

2.4.3.6 Maximum Diversification Assuming that all stocks have the same return is unlikely to be true. (Choueifaty and Coignard 2008) propose an optimization method to maximize the diversification of a portfolio by instead assuming that returns are proportional to the risks of each stock. With R_f being the risk free rate, and the expected return given by $E(R_i)$, this relationship is given by:

$$E(R_i) - R_f = \gamma\sigma_i \tag{9}$$

where $\gamma > 0$ is the proportion of return per unit of risk.

The second major contribution of (Choueifaty and Coignard 2008) is the quantifiable measure of diversification known as the Diversification Ratio (DR).

$$DR = \frac{\sum\limits_{i=1}^{N} w_i\sigma_i}{\sigma_p} \tag{10}$$

DR is the ratio of the weighted sum of the individual stock volatilities ($\sum_{i=1}^{N} w_i\sigma_i$) to the fully correlated and weighted sum of volatilities given by the portfolio volatility (σ_p).

Solving for the vector of weights that maximizes the diversification ratio is given by:

$$\underset{w}{max}\left(\frac{w'\sigma}{\sqrt{w'\Sigma w}}\right) \text{subject to } \begin{cases} \sum_{i=1}^{N} w_i = 1 \\ 0 \leq w_i \leq max\% \end{cases} \tag{11}$$

For comparison purposes relative to the minimum variance approach, I use maximum position sizes of 1% and 5%. As can be seen in Table 17, even at these lower position size limits, the approach can lead to a small selection of securities for the portfolio.

The original work of (Choueifaty and Coignard 2008) was based on stocks in the S&P500 index from 1990 to 2008 and used 250 days of data to construct the parameter estimates required for the covariance matrix. Position limits were set at 10% but were additionally limited to month end weights which complied with UCITS III[4] rules. This restriction included a maximum limit of 40% on the sum of individual stock weights which exceed 5%. Even with these additional restrictions, setting position limits at this level often leads to very concentrated portfolios.

My simplified yet comparable approach (which does not include the UCITS restrictions) is calculated across the time frame of 1968 to 2017, from a universe of 1,000 stocks. This method on average carried an effective number of stocks of just 24. Using this method but with 10% position limits, I observe that although the methodology maximizes the ex-ante DR, the ex-post DR is only slightly above that of the market cap weighted method and on average below all of the other methods examined in this study. I do not include the 10% max position size version in my analysis but do note the effects of a high concentration risk as will be shown in my research when combining the steps of capital allocation with perfect stock selection.

[4] `https://www.esma.europa.eu/databases-library/interactive-single-rulebook/ucits`

2.5 Properties of the Returns

The wide array of data presented in Table 1 gives an opportunity to describe some of the notable differences observed in the characteristics of the statistics. Each of the allocation methods can be compared against the base method of market capitalization and against the other alternative capital allocation methods.

2.5.1 Returns:

The metric of return is a measurement of the payment that investors earn over a period of time, presumably for the assumption of some risk. Any estimate of this metric will be imperfect due to issues such as auto-correlation, mean reversion at long horizons, etc. In addition, observations of stock returns over various horizons depict a distribution which is not normal. Therefore the calculation of an unbiased estimate of return is problematic.

In an early paper, (Blume 1974) documented the biases that exist with both arithmetic and geometric returns. In general, arithmetic returns are biased to the upside while geometric returns are biased to the downside. The author developed four unbiased estimators of return and showed their efficiency. In a subsequent paper, (Jacquier, Kane, and Marcus 2003) showed that the unbiased estimate is a weighted average between arithmetic and geometric returns. For short periods, arithmetic is heavily weighted and as the horizon return begins to approach the estimation period data, geometric is most heavily weighted.

Advanced methods for correcting for these biases and designing an efficient estimator was proposed by (Jacquier 2005).

An arithmetic return estimate is constructed as the simple mean of the observed single

Table 1: Summary statistics for various capital allocation methods, 1968 - 2017

	MC	**EW**	**IV**	**ER**	**MV1%**	**MV5%**	**MD1%**	**MD5%**
Geom Ret %	10.2	11.7	12.2	12.0	11.6	11.2	11.4	10.9
Arith Ret %	11.4	13.3	13.5	13.3	12.4	11.9	12.4	12.0
Skewness	-0.39	-0.35	-0.33	-0.38	-0.40	-0.66	-0.60	-0.65
Excess Kurtosis	1.77	2.16	2.46	2.54	3.63	2.72	2.44	2.47
Ann Vol %	15.1	17.0	15.5	15.2	11.9	11.3	13.8	14.3
Down Vol %	9.6	10.7	9.6	9.4	7.3	7.0	8.8	9.3
Track Error %		5.4	4.9	4.8	7.2	8.0	6.3	7.6
Max DD %	-50.5	-52.2	-49.5	-48.9	-38.6	-36.8	-42.7	-45.5
DaR 5%	-33.7	-27.6	-26.2	-25.5	-19.9	-19.1	-21.5	-22.3
CDaR 5%	-45.1	-42.9	-44.0	-43.5	-35.1	-33.7	-34.8	-37.7
Sharpe Ratio	0.44	0.50	0.56	0.56	0.64	0.63	0.55	0.50
Modified Sharpe	0.17	0.19	0.21	0.20	0.22	0.21	0.20	0.18
Sortino Ratio	0.70	0.80	0.92	0.91	1.06	1.03	0.88	0.79
Info Ratio		1.88	2.06	2.1	1.4	1.3	1.6	1.3
Ret / DaR	0.32	0.45	0.48	0.49	0.59	0.59	0.55	0.51
Effect # Stocks	138	1,000	877	831	113	52	113	44
Divers Ratio	2.04	2.16	2.13	2.21	2.33	2.35	2.65	2.80
Avg Turnover %	8.8	13.0	12.5	13.9	36.8	43.2	40.3	47.0
Active Share		47.1	49.2	52.1	84.8	90.1	88.0	93.6

MC: 'Market Capitalization' holds all 1,000 stocks and weights the holdings annually based on the relative market capitalization of each stock to the total market capitalization of all stocks on that date. **EW:** 'Equal Weight' holds all 1,000 stocks and weights each stock equally. **IV:** 'Inverse Volatility' gives more weight to stocks with lower volatility and less weight to stocks with higher volatility as measured monthly over the previous five year. **ER:** 'Equal Risk' holds all 1,000 stocks, weights solved by optimization method whereby the marginal contribution of each stock to total portfolio risk is equalized. **MV1%/5%:** 'Minimum Variance' holds a subset of the available stocks where weighting is based on an optimization which finds the portfolio with the minimum total variance, subject to maximum position sizes of 1% or 5%. **MD1%/5%:** 'Maximum Diversification' holds a subset of stocks each year with weights based on an optimization which maximizes the ex-ante diversification ratio (DR) of the portfolio using volatility and correlation data monthly from the previous five years. Position sizes are limited to 1%, 5%. **Statistical calculations contained in this table are explained in the appendix of this chapter.**

period returns. Geometric returns are calculated as the single period return that would allow an investor to earn the total compounded return observed over the sample period. For the analysis in this dissertation, both metrics are calculated and any biases that would appear from using these estimates for forecasting future returns are inherent symmetrically across each of the allocation methods. With a more thorough basis for understanding return estimates, the focus now turns to a comparative look across the return statistics of each allocation method.

In general we see that both geometric and arithmetic returns for the market capitalization approach are the lowest while those of inverse volatility and equal risk are the highest. All alternative allocation methods outperform the market capitalization approach for both geometric and arithmetic return measures for the 50 years of this study. For the purpose of this research I focus on the arithmetic returns rather than geometric returns. Each has particular merits and restrictions.

2.5.2 Return Distribution Normality:

Distributions with negative skew have more observations in the left tails than would be implied by the normal distribution. This is a bad characteristic for investors and should require a higher return premium as compensation for this risk. Skewness for the market capitalization and the three allocation methods (EW, IV, ER) that contain all 1,000 stocks are relatively well behaved. But for the minimum variance and the maximum diversification approaches, negative skewness is more pronounced with the worst observations being at the 5% maximum position size for each of these optimization methods.

Kurtosis measures the weight in both tails together. A normal distribution has a kurtosis of 3. This analysis shows excess kurtosis or the amount in excess of a normal distribution. The market capitalization approach has the lowest observed excess kurtosis at 1.77. By contrast, the minimum variance methods shows the highest levels of excess kurtosis.

The combination of excess kurtosis and negative skew are especially bad and would naturally be expected to warrant a higher return premium as compensation. The worst combinations of excess kurtosis and skew are observed in the minimum variance methods. While the additional premium required for this combination is not seen simply in the returns of these approaches, they do show the highest ratios of reward vs. risk such as Sharpe, Sortino, and Return/DaR which is consistent with this expectations. The Modified Sharpe ratio (explained below) accounts for these elements of non-normality in the distributions of returns.

2.5.3 Volatility:

The volatility for the market cap weighted portfolio sits in the middle of all of the other approaches. Equal weight has the highest volatility which may be due to the size bias that would naturally exist in a method which spreads the allocation weights equally relative to an approach which is highly concentrated in the largest capitalization names. As would be expected, the minimum variance approaches show the lowest realized volatility but this is clearly a function of the metric that is optimized during the construction of the portfolio.

2.5.4 Drawdown:

The are three metrics to examine in this area. All values are calculated on a monthly basis and indicate the loss from a previous peak in value.

Maximum drawdown is the single largest loss event experienced over the entire history. While most of the other allocation methods show improvements in this area, the equal weight method has larger losses during the worst single event. The greater size of drawdowns seen in the market cap method are likely functions of the effective concentration of these portfolios into a very few stocks. The large drawdowns from the equal weight method may be explained by the increased exposure to smaller cap securities.

Drawdown at Risk (DaR 5%) examines the distribution of all drawdown events and denotes the size of the 5% worst event. It is important to note that each complete event (peak to trough to peak) is treated as a single observation for the distribution. Once we move beyond examining just the single largest drawdown event, the DaR 5% tells a different story. Every other allocation method has better metrics with improvements in the range of 6.1% to 14.6% better than MC. During a relatively normal sized drawdown event, any method other than the market capitalization allocation method performed better across this time frame.

The final metric is the **Conditional Drawdown at Risk** (CDaR 5%) which examines the average size of the drawdown which is in excess of the DaR 5%. CDaR 5% calculates the mean loss that occurs when the loss is within the left tail. Maximum drawdown is the single worst observation in the distribution and it will have a direct impact on the CDaR calculation. It is interesting to note that despite the equal weight method having larger absolute maximum drawdowns, the market capitalization method still exhibits the highest CDaR 5%.

2.5.5 Reward vs. Risk:

For both the Sharpe ratio and Sortino ratio, the market capitalization method shows the lowest mean statistics. As noted above, the minimum variance (1% max position & 5% max position) methods appear to be clear winners for nominal risk adjusted returns, but this will be challenged with robustness testing. These higher risk adjusted returns may be the required compensation for the combination of negative skew and excess kurtosis that these strategies exhibit as the Modified Sharpe Ratio (below) for each of these methods are in line with the other allocation methods. All other allocation methods appear to outperform the market capitalization method, but this will be tested.

The **Sharpe Ratio** was originally proposed by (Sharpe 1966) as the reward to variability ratio. In it's original form, the equation given was:

$$R/V = \frac{A_i - p}{V_i} \tag{12}$$

with A_i being the return of the investment, p being the risk free rate, and V_i being the standard deviation of the returns on investment i. A_i was described by Sharpe as the 'average annual rate of return'. Unfortunately, there is no single statistical measure known as 'average' and so I default to the most simple approach which is the arithmetic average of returns. Practitioners sometimes use geometric averages.

The Sharpe ratio is a common measure of risk adjusted returns, however as the denominator is simply the standard deviation of the return series, it suffers from being indiscriminate relative to distributions with high levels of skew and excess kurtosis which are generally disliked by investors.

The **Modified Sharpe** was originally suggested by (Favre and Galeano 2002) and then by (Gregoriou and Gueyie 2009). The Modified Sharpe ratio replaces the denominator of the original Sharpe ratio (standard deviation of returns) with a modified Value at Risk (mVaR) calculation which accounts for the extreme observations more fully by incorporating skew and excess kurtosis. With a standard parametric VaR measure, a z-score is used as a scalar on the standard deviation of the observations. The mVaR calculation uses a Cornish Fisher expansion to adjust the z-score for the non-normality of the distribution before scaling the standard deviation observation.

The mVaR ratios are not directly comparable to standard Sharpe ratios. They are however more comparable to each other as each has been properly adjusted for the first four moments of their distributions.

A case in point can be made by looking at the ratios for the IV and ER allocation methods. Standard Sharpe ratios are both 0.52 while the modified Sharpe ratios are different. As ER has slightly more negative skew and a higher excess kurtosis, the modified Sharpe is slightly lower with 0.21 (IV) and 0.20 (ER). In cases where distributions have extreme third and fourth moments, the modified Sharpe is key to adjusting these ratios and leveling the comparisons for risk adjusted returns.

The **Information Ratio** is highest with the inverse volatility and equal risk contribution methods which is a function of the low tracking error that these method exhibit. As tracking error is the denominator of the information ratio, the methods with the highest tracking errors (minimum variance and maximum diversification) have the lowest information ratios. These two methods have the largest differences in construction as will be shown with Active Share and while there are risk adjusted performance advantages to straying from the benchmark, the excess tracking risk and information ratios are negatively impacted by these approaches.

2.5.6 Diversification:

There are several ways to consider diversification. The most simplistic approach is simply the number of stocks in the portfolio. The market capitalization, equal weight, inverse volatility, and equal risk methods the portfolios hold all 1,000 stocks that are available on each date. The minimum variance and maximum diversification portfolios hold a smaller subset of stocks due to the optimization routine preferring some stocks to carry zero weight.

Another measure of diversification to consider is the **Effective Number of Stocks** which is a function of the allocation weights. This metric is calculated as the inverse of the Herfindahl-Hirschmann Index (HHI)[5] which is a measure of concentration.

$$HHI = \sum_{i=1}^{N} w_i^2 \tag{13}$$

The inverse of the HHI is the Effective Number of Stocks. Large weights to a few stocks will make this ratio smaller while low weights to many stocks will make this ratio larger.

Despite holding 1,000 stocks, the market capitalization method is only effectively influenced by 138 stocks. This is in stark contrast to the equal weight method which holds the same exact stocks but in equal proportion so that the effective number of stocks held is exactly 1,000. Similarly, the inverse volatility and the equal risk portfolios have very high effective number of stocks and would appear to be more diversified than the market capitalization method, even though they all own the same set of stocks each year. With a smaller set of stocks held and relatively large concentration weights, the minimum variance and maximum diversification methods display a metric for effective number of stocks that is less than the market capitalization method.

A third method to consider is the **Diversification Ratio** which is a function of weights, volatilities, and the correlations between securities. While the maximum diversification method seeks to maximize this ratio ex-ante, the ratio presented in Table 17 on page 237 is the ex-post result. Again we see the market capitalization method being the least diversified and all other allocation methods are superior in this ratio. As noted previously but not included in our analysis, while the maximum diversification (10% max position) is the least constrained version which attempts to maximize this ratio, it barely shows any difference relative to the most concentrated method which is the market capitalization. There appears to be a wide spread between the optimization of

[5] `https://corporatefinanceinstitute.com/resources/knowledge/finance/herfindahl-hirschman-index-hhi/`

this ratio ex-ante and the ex-post resulting DR when position sizes are allowed to range quite high.

2.6 Turnover:

The market capitalization method has the lowest turnover at 8.8% per year on average. This is unsurprising as the weights of stocks in the portfolio float with changes in value and are only influenced by names being added or deleted from the universe once per year. The second lowest set of turnover is seen with the three methods that hold the same set of 1,000 stocks but have their weights changed independently from the value of the stocks. These are the equal weight, inverse volatility, and equal risk contribution methods. Finally, the minimum variance and maximum diversification methods change the subset of stocks held and solve for a new set of weights each period which drives their average annual turnover much higher than the other methods.

A final metric shown was put forth by (Cremers and Petajisto 2009). **Active Share (AS)** uses the current holdings to measure the 'active bets' that a manager is making relative to the benchmark. The authors explain the measure as 100 minus the overlapping position weights. AS can increase as stocks are either included or excluded and as the weights of stocks vary relative to their weights in the benchmark.

"Active management" tends to refer to managers who engage in security selection. However, according to ranges of AS set by the authors, all of the alternative allocation methods in this study show AS scores in the range of active management even though none of them are engaging in intentional stock selection. Equal weight, inverse volatility, and equal risk contribution hold all of the same stocks as the benchmark but only vary their weights. Minimum variance and maximum diversification hold a subset of stocks and so the higher AS measure are a function of both of these variables. Minimum variance and maximum diversification methods show AS scores in the range of 85-95.

2.6.1 Turnover and Transactions Costs

While all of the capital allocation methods chosen for my panel have higher returns, they also have higher turnover ratios. This means that in order to maintain the allocation method, there is more trading each year in all of these methods than there is in the market capitalization method. Increased trading may indicate higher costs and thereby lower returns. Adjusting these return streams for potential costs would seem to be a logical way to try to compare strategies on a more even basis.

The market capitalization method still has some turnover. Each year, some stocks decrease in value while others increase. Some firms merge together or there are spin-offs that change the composition. All of these events have cause an annual one-way turnover in the market capitalization method of 8.8% per year. In order to make an adjustment for trading costs, the comparison should only use the additional turnover of the alternative strategies above the reference method, market capitalization weighting. The other allocation methods range from a low of 12.5% (IV) to a high of 47% (MD5%).

It seems logical that more trading would diminish the realized returns of these strategies over time. But adjusting the returns downward for assumed transactions costs is fraught with potential errors which are described below.

- Market Structure - over the 50 year period in examination, the market structure has greatly changed. Dealers have entered and left the market, hedge funds and high speed traders have become larger players, and retail trading has increased dramatically. Also, market regulation has changed for each of these participants. As a market is just a collection of traders, this shifting market structure has had a meaningful change on the liquidity and costs for all market participants.

- Commission Rates - the explicit costs of trading have fallen dramatically over the last 50 years. In addition, the commission rates paid by traders of various size can vary greatly. Large traders pay very small explicit commissions, while small retail traders have paid much higher prices in the past.

- Bid/Ask Spreads - crossing the bid/ask spread is an implied cost of trading. As the market structure has changed over the past 50 years, the bid/ask spreads in stocks has generally decreased. In addition, the spread is not the same for all securities and can even change for a single security across time. Making a single assumption for all stocks across the full 50 year time frame would create potential errors in the return assumptions.

- Trade Size - trading very small or very large order sizes can have an impact on the total cost of the trade. The size of the trade and sophistication of the trader can allow the order to access additional liquidity in dark pools or find off market liquidity with dealers through a 'Request for Quote' (RFQ) process.

- Constellation of Liquidity - it is not just the market in the single common stock that impacts the total cost of the trade. The full constellation of liquidity must be considered which includes stocks with high correlations, listed options markets in the common stock, capital structure trades vs. corporate bonds in the same issuer, and futures markets. All of these factors contribute to the total liquidity of a stock and each of these factors is constantly in flux which makes assumptions about liquidity and total cost of trade very difficult to estimate.

- Timing - the speed at which the trade must occur can also impact the total cost of trading. If a trade must happen instantaneously the cost naturally rises. However, if a trade can be filtered into the market over several minutes, hours, or days, then the total cost can be expected to decline.

- Frequency - high speed algorithms available today make the management of a trade order very specialized. The ability to slice orders into tiny pieces within the depth of the top of the market book can help to decrease the overall cost. Also, the ability to float the order on the bid or offer side without crossing the spread helps to provide liquidity to the market and can actually make the net cost of the trade negative. Intelligent trade order management can create small profits which implies that the strategies with higher turnover could actually be sources of profit for the strategies rather than the initial expectations of decreasing returns.

With all of these potential factors impacting the total cost of trading, it is easy to see that a single set of assumptions for each of these items across the full 50 years would be impossible. Adjusting returns downward for trading costs is not a feasible way to approach this problem.

The standard approach within the literature is to look at the problem from the other side and extract the net costs of the additional trading that would be required in order to make the two strategy returns equal. Then this total cost can be considered whether it appears reasonable or not. This general approach has been used in many papers including (Clare, Motson, and Thomas 2013).

In the table above, the per trade marginal cost is computed which would make the alternative allocation method returns equal to the market capitalization method. The top section of the table shows the annualized returns of each strategy, and the marginal returns of the panel of allocation methods relative to the market capitalization method. The second section displays the annualized turnover of each strategy, and the marginal turnover of the panel of allocation methods relative to the market capitalization method.

Table 2: Marginal Trading Cost Analysis

	MC	EW	IV	ER	MV1	MV5	MD1	MD5
Return %	10.2	11.7	12.2	12.0	11.6	11.2	11.4	10.9
Marginal Return %		1.5	2.0	1.8	1.4	1.0	1.2	0.7
Turnover %	8.8	13.0	12.5	13.9	36.8	43.2	40.3	47.0%
Marginal Turnover %		4.2	3.7	5.1	28.0	34.4	31.5	38.2
Implied Cost %		35.7	54.1	35.3	5.0	2.9	3.8	1.8

The table above shows the computed total cost of trading required in order to reduce the annualized returns of the panel of allocation methods downward to equal the returns of the market capitalization method. The method of computing the implied costs is to divide the marginal return percentage by the marginal turnover percentage.

The implied costs percentage is then computed as:

$$Implied\, Cost\, \% = \frac{Marginal\, Return\, \%}{Marginal\, Turnover\, \%}$$

The first three strategies in the panel (EW, IV, ER) all have the highest marginal returns and the lowest marginal turnover. This gives implied costs of 35.3% to 54.1%. Clearly these costs of trading appear far in excess of the expected costs for large capitalization US stock trading. However, the MV and MD strategies have higher marginal turnover and lower marginal annual returns. This has the effect of decreasing the implied costs into the single digit percentage levels. These range from a low of 1.8% to a high of 5.0%. While it is impossible to know the liquidity of the exact stocks traded on the dates traded, an unskilled trader may begin to encounter actual costs that approach these levels.

My analysis of transactions costs is inline with the observations from existing literature, which is validation of the methods that I have used to create these portfolios. And while this style of analysis can call into question the ability to achieve the returns of some of these strategies, this argument has been made in other literature and is not the focus of my questions and analysis.

Acknowledgement of increased turnover and potential transactions costs is important to preserve the practical nature of this thesis. However, the focus of my research is in two parts. First, I show how the ability to choose stocks well or poorly impacts these various allocation methods. Second, I show how selecting stocks based upon returns has less impact than selecting stocks based on correlations. The relationships that I will show in the research are preserved in the presence of transactions cost differences.

2.6.2 Factor Exposures

2.6.2.1 Fama-French 3 Factor Model (FF3): This asset pricing model, developed by (Fama and French 1992), extends the original single factor capital asset pricing model (CAPM) by including exposure sensitivity to firm size and valuation ratios. These additional factors have been shown to provide greater explanation for individual equity returns and help to increase the R^2 of the regression.

Exposures (betas) for the monthly returns of each portfolio allocation method are shown in Table 18 relative to the factors of market, size, and value as provided by Kenneth French[6]. Calculations are based on a multiple regression analysis and calculated in Python using the Statsmodels[7] package which is documented by (Seabold and Perktold 2010).

It is important to note that although I have constructed a market capitalization weighted portfolio as the benchmark for our analysis, our methodology differs slightly from that used for the creation of the factor returns. Stock inclusion may be slightly different, and

[6] `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`
[7] `https://www.statsmodels.org/stable/index.html`

Table 3: Fama-French Three Factor Model, 1968 - 2017

	MC	EW	IV	ER	MV1	MV5	MD1	MD5
Intercept	0.00%	-0.02%	0.00%	0.00%	0.12%	0.12%	0.07%	0.04%
t-stat	-0.07	-0.47	0.67	0.85	1.94	1.78	1.12	0.54
Market	0.99	1.04	0.96	0.94	0.71	0.65	0.81	0.79
t-stat	216.87	102.22	90.40	91.57	48.81	41.54	55.88	43.67
SMB	-0.12	0.27	0.17	0.17	0.04	0.05	0.18	0.22
t-stat	-17.94	18.75	10.93	11.60	1.88	2.15	8.47	8.36
HML	0.06	0.27	0.33	0.30	0.27	0.25	0.21	0.20
t-stat	8.16	17.62	20.05	18.90	12.19	10.60	9.69	7.36
R^2	0.99	0.96	0.94	0.94	0.81	0.76	0.86	0.80

Monthly excess returns above the risk free rate for each allocation method are regressed against the following set of standard factors as collected from: `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/index.html`. The website contains full information on the formation of the factor return sets. **Market:** excess return above the 1 month US Treasury Bill of the value weighted return of all CRSP firms incorporated in the US and listed on NYSE, AMEX, or NASDAQ exchanges with a share code of 10 or 11. **SMB:** Small Minus Big - the average return on the three small portfolios minus the return of the three big portfolios. **HML:** High Minus Low - the average return on the two value portfolios minus the average return on the two growth portfolios. Capital allocation methods: MC - market capitalization, EW - equal weight, IV - inverse volatility, ER - equal risk, MV1 - minimum variance 1% max position size, MV5 -minimum variance 5% max position size, MD1 - maximum diversification 1% max position size, MD5 - maximum diversification 5% max position size

we rebalance annually while the factor return portfolios are rebalanced quarterly.

- **Intercept:** The intercept in the equation is the additional performance not explained by the combination of factor exposures solved in the multiple regression. Using a factor lens, the intercept is the unexplained performance alpha (α). Given the performance improvements that are seen in the gross return numbers for the alternative allocation methods, the difference in the intercepts relative to our version of the market cap portfolio are quite small as most of the performance is explained by combining market returns with a few additional factors.

The t-statistic is shown for each value of the intercept. Our null hypothesis is that the true alpha of any of these allocation methods is zero. Using a two tailed test at the 95% confidence level, this would make the critical value equal to 1.96 which would cause us to reject the alphas for all of the methods although the alpha for MV1 is very close to being significant.

Minimum variance portfolios are in several characteristic ways the most different than other allocation methods and have the largest intercepts although we have to reject these values at the 95% confidence level and assume that they are possibly zero.

- **Market:** The market exposure factor is the excess return of a value weighted portfolio of stocks above the risk free rate. For the three methods which hold all 1,000 stocks, the market exposure is quite high and very near that of the market capitalization method in our panel. As minimum variance and maximum diversification hold fewer stocks and have dramatically different weightings, the portion of returns attributable to the overall market are significantly lower. All of the t-statistics are significant for the market coefficients.

- **Size (SMB):** The size factor captures the return of small versus big stocks (market capitalization) within a portfolio. A positive value suggests more exposure to smaller capitalization stocks than in the market portfolio. As expected, all of the allocation methods examined in this study show more exposure to smaller capitalization stocks. This is by design as the market cap weighted portfolio concentrates its exposure in the largest stocks available. Of the methods that hold all 1,000 stocks each year, equal weight has the most exposure to smaller cap stocks. All of the t-stats are significant with the exception of MV1 which is close to the threshold.

- **Value (HML):** The value factor captures the excess return of cheap stocks relative to expensive stocks. This is measured through the book to market ratio. The full panel of allocation methodologies display additional exposure to the value factor in the range of 0.20 - 0.33. This is reasonably significant with all of the t-stats above the threshold limit of 1.96.

- **R Squared:** The allocation methods which hold the full 1,000 stocks (equal weight, inverse volatility, equal risk) are very well explained by the factors in the model with a minimum R^2 of 0.94. While the factor model explains less of the optimized methods (minimum variance, maximum diversification) returns, the R^2 is still relatively high. It is noticeable that the more concentrated strategies (max position sizes of 5%) have lower explanatory power by the factor model.

2.7 The Hierarchy of Sharpe Ratios

2.7.1 Why Focus on Sharpe Ratio?

Of all the statistical characteristics which have been described for our panel of capital allocation methods, my research focuses on the risk adjusted returns as given by the Sharpe Ratio.

(Markowitz 1952) set the stage for this risk-adjusted focus before (Sharpe 1966) proposed his famous ratio.

> *"The hypothesis (or maxim) that the investor does (or should) maximize discounted return must be rejected. If we ignore market imperfections the foregoing rule never implies that there is a diversified portfolio which is preferable to all non-diversified portfolios. Diversification is both observed and sensible; a rule of behavior which does not imply the superiority of diversification must be rejected both as a hypothesis and as a maxim."*

Sharpe's work was an expansion of the earlier work by (Treynor 1965) who first incorporated the volatility of a fund's return into a method for predicting future performance. Sharpe emphasized the point made by Markowitz earlier:

> *"The key element in the portfolio analyst's view of the world is his emphasis on both expected return and risk."*

In (Modigliani and Modigliani 1997), the authors further explain Markowitz's assertion that returns are "an incomplete measure of the performance of a portfolio because it ignores risk." An alternative measure of 'Risk Adjusted Performance' (RAP) is introduced which is intended to be easier for investors to understand. The method adjusts all portfolio returns to match the market risk of the benchmark and allow for easy comparison of returns which they assert is grounded in finance theory but easier to understand.

Since the publication of (Sharpe 1966), many authors have expanded on the work with more complicated versions which account for higher moments, non-normality, etc. Additionally, (Hwang, Xu, and In 2018) use a focus on tail risk to explain the performance differences between naïve and optimal allocation. These two examples show the growing complexity of the literature with regard to risk adjusted performance.

Sharpe ratios have several assumptions built in with regard to the distribution. As financial time series data does not adhere to a normal distribution, the resulting statistics of Sharpe ratios can have peculiarities. In (Lo 2002), the author explains how the Sharpe ratio can not be expanded with the square root of time assumption, and how annual Sharpe ratios can be overstated to a large degree because of serial correlation. All of this can negatively impact the ranking of fund returns making them less valid in a hierarchy.

The modified Sharpe ratio has a fair degree of complexity where the denominator is a Cornish-Fisher expansion of a normal z-score incorporating kurtosis and skewness. While this appears to be a more complete way to represent the true risk or a distribution, it creates other issues. First, it is not easily interpreted by investors. Second, it is not directly comparable with the original Sharpe ratio measure. However, for the purpose of comparison, Table 17 shows the calculation of the Modified Sharpe Ratio (mVaR) for each allocation method and the Appendix contains the mathematical description.

While there are benefits to more advanced measures of risk adjusted performance, the downside is that they are not easily understood. From a practitioner's perspective the focus will be on the most commonly quoted of risk-adjusted ratios, the Sharpe Ratio[8], with the knowledge that its simplicity masks some of the true complexity of risk adjusted performance.

In addition to simplicity of the calculation, and market acceptance of the measure, the authors of (Eling and Schuhmacher 2007) show that the choice of measure rarely matters when ranking the risk adjusted performance of funds. Using the Sharpe ratio and 12 other performance measures, a study of 2,763 hedge fund returns show that when using any of these measures, the result is a nearly identical rank ordering.

It is widely documented that portfolio returns are not normally distributed and have serial correlation. All of this can create issues with fairly representing the risk of a portfolio, the basis for all risk adjusted performance measures. But despite the proliferation of methods for judging risk adjusted performance, and the sophistication of alternative measures, very little difference in the rankings is observed. Additionally, the investment marketplace has widely accepted Sharpe ratio as the default risk adjusted measure of performance. For this combination of reasons, and given the practical nature of my ex-

[8]Sharpe originally named this the "reward to variability ratio" but authors began to refer to it as the Sharpe ratio and this has been accepted in the literature.

amination, a focus on the original Sharpe ratio was selected.

2.7.2 Hierarchy - Initial Results

A scan of the Sharpe Ratios calculated for each allocation method across the 50 year time horizon yields a hierarchy of risk adjusted performance. Initial results confirm prior literature showing that all systematic allocation method Sharpe ratios are higher than that of the market capitalization method.

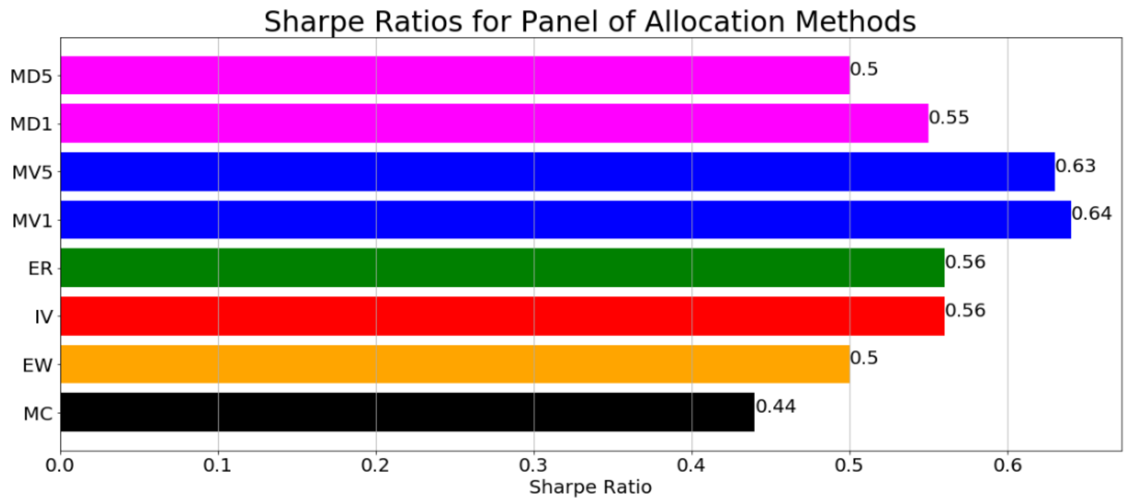

Figure 3: Hierarchy of Sharpe Ratios
The figure above shows the average annual Sharpe ratio of each allocation methodology estimated using monthly returns over the 50 year period of 1968 - 2017.

It is important to note however that while these results are generally in line with prior literature observations, they are calculated over a longer time horizon. The works of both (Clare, Motson, and Thomas 2013) and (Chow et al. 2011) both include the financial crisis of 2008 and are extended to 2011. This thesis has the same starting date of 1968, but is extended beyond 2011 to include 2017. During that additional six year period, the extremely easy monetary policy of the central banks around the world led to the continuation of a great bull market in equity returns.

It may be that the hierarchy of Sharpe ratios has changed over shorter periods of time. Researchers have used shorter time periods to show the instability of a single observation. Often this is done in decades, given such as 1970 - 1979. These are arbitrary starting and ending periods. While not the focus of this research, if an analysis of the characteristics of a strategy were to be the focus of study, then it may be that a definition of economic epochs could be used to show this shift in hierarchy. This may take the form of a bull market vs. bear market period, or periods of easy monetary policy vs. tight monetary policy as examples. These epochal definitions may be designed in a nearly unlimited number of ways.

However, as the time period expands, the estimation of the hierarchy converges to the true hierarchy. A 50 year period now available for examination gives more probability to observing a risk adjusted performance hierarchy that represents the true characteristics of these allocation methods.

These initial results require further examination as they are simple point estimates for a single period of time. To understand this hierarchy more fully, we need to examine the robustness of each pair of Sharpe ratios. The matrix of differences between each allocation method is shown in Table 4. The values in the matrix represent the gross difference between each of the pairs of allocation methods across the entire 50 year time period.

2.8 Robustness Test - Description

The observed Sharpe ratios form a hierarchy. But these measurements are ex-post point estimates and as such can not be relied upon with precision. It cannot be said for certain that the Sharpe ratio estimate for inverse volatility (0.50) is greater than the estimate for equal weight (0.44). In order to properly assess that the observed (but es-

	EW	IV	ER	MV1	MV5	MD1	MD5
MC	-0.06	-0.12	-0.12	-0.20	-0.19	-0.11	-0.06
EW		-0.06	-0.06	-0.14	-0.13	-0.05	-0.05
IV			0.00	-0.08	-0.07	0.01	0.06
ER				-0.08	-0.07	0.01	0.06
MV1					0.01	0.09	0.14
MV5						0.08	0.13
MD1							0.05

Table 4: Sharpe Ratio Differences

The table above shows the observed difference between two Sharpe ratios over the 50 year test period. The number is constructed as the row header minus the column header. For the entry in the upper left of the table, this would be the observed Sharpe ratio of the MC method less the observed Sharpe ratio of the EW method. The number is negative as EW has a higher observed Sharpe ratio than the MC method.

timated) hierarchy is significant, we must robustly test for the differences between all Sharpe ratios.

The original method for performance hypothesis testing was given by (Jobson and Korkie 1981) who developed an asymptotic distribution of the estimator. This test was later corrected by (Memmel 2003). However, in (Ledoit and Wolf 2008), these original tests are shown to be invalid for distributions that have excess kurtosis or are based on a time series of returns due to serial correlation. Kurtosis is a common problem with financial market returns, while serial correlation is less problematic for liquid markets and more of an issue for hedge funds and illiquid investments. Their research proposed the use of a robust inference method which constructs a studentized time series bootstrap confidence interval for the difference between two Sharpe ratios. The difference of the two ratios is considered significant if zero is not contained in the interval calculated.

To complete these tests, my analysis has benefited from the open source code of (Ledoit and Wolf 2008) which is available in both MATLAB and R format.[9]

The algorithm to test the difference between Sharpe ratios begins by fitting a semi-parametric model to the observed data. The next step creates a set of bootstrapped time series by using various block lengths. In a normal bootstrap procedure, returns are drawn one at a time with replacement from a set of observed data. To overcome the issue with serial correlation, returns are selected in blocks of sequential dates chosen at random. For monthly returns with block size = 2, adjacent months are chosen together to form the bootstrap of the required length. In the calibration phase of this procedure, various block sizes are tested. The block size that produces a distribution which most closely matches that of the semi-parametric model within a set confidence level is the optimal block size.[10]

While the calibration of the optimal block size for each set of Sharpe ratio differences converges to its own answer, for a consistent comparison I have used a block size of 5 as the average observed in the calibration stage. It is important to note that the mean difference in Sharpe ratios ($\hat{\Delta}$) does not change with different block sizes and that the p-value of the difference is relatively insensitive to the choice of block size within a reasonable range. In addition, in Tables 5 and 6, I denote both the 95% and 90% confidence levels to adequately account for the slight changes in p-values that may occur from selecting a block size other than the most optimal for each Sharpe ratio pair.

While Table 4 is shown in terms of annualized Sharpe ratios, the code for the robustness test of (Ledoit and Wolf 2008) uses monthly Sharpe ratios. As explained in (Lo 2002), monthly Sharpe ratios do not convert to annual Sharpe ratios through the simple multiplication by $\sqrt{12}$ except under very specific circumstances. This difference in

[9] `https://www.econ.uzh.ch/en/people/faculty/wolf/publications.html#9`

[10] The calculation of optimal block size is computationally intensive. On an i7 processor @ 3.4 GHz with 32gb RAM and running the MATLAB code, it can easily take 5-10 hours to converge to the answer depending on the number of block sizes submitted for testing.

monthly vs. annual Sharpe ratios does not impact the analysis or its conclusions, only the presentation of levels between Table 4 and Table 5.

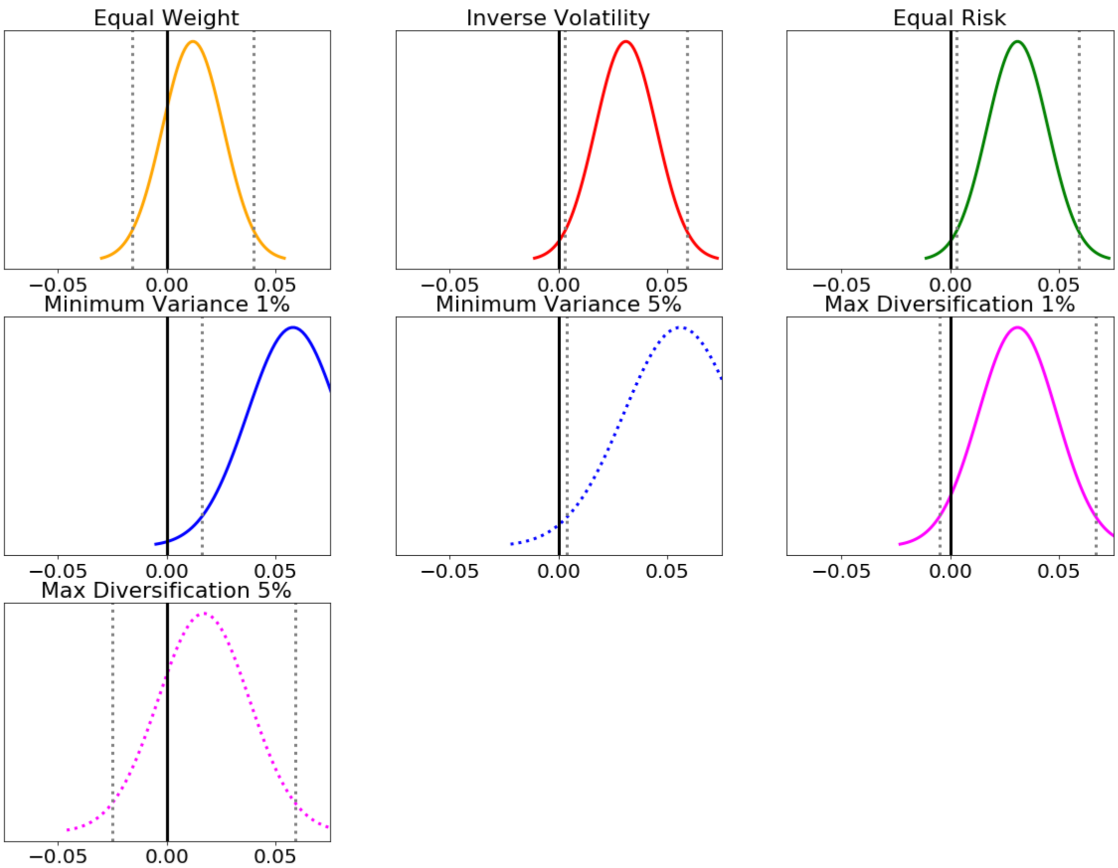

Figure 4: Statistical Significance of Sharpe Ratio Differences
The figure above shows the distribution of monthly Sharpe ratio differences for each allocation methodology relative to the market cap allocation method. The dashed vertical lines denote a two-tailed 95% confidence interval. When zero is contained within this confidence band, then the null hypothesis can not be rejected and it must be assumed that the true difference in Sharpe ratios is equal to zero.

2.8.1 Robustness Test - Example Results

Each of the visuals in Figure 4 shows the difference between the market capitalization method and one of the other allocation methods in the panel. The distributions are calculated from the robust bootstrap testing. Each distribution includes the standard deviation of the results of a simulation of 5,000 bootstrap tests on the difference between two Sharpe ratios.

For example, in Figure 4, the chart in the upper left is the distribution of monthly Sharpe ratio differences between the market cap method and the equal weight method. The mean of the distribution is above the zero line which is consistent with our results from Table 4. However, this visual gives more information which extends beyond a simple ex-post point estimate.

The mean and the standard deviation from the bootstrap simulations form the orange distribution curve. The two vertical dashed lines depict +/- 1.96σ for a two tailed 95% confidence band. Since the zero value is contained within the 95% confidence band, I can not reject the null hypothesis that the difference between Sharpe ratios is indistinguishable from zero. While it first appeared that the equal weight allocation method has a superior Sharpe ratio, this cannot be accepted as our test shows a lack of statistical significance.

In Figure 4, the left most entry on the second row shows the same information for the difference between MC and the MV1 (blue curve). As zero is not contained within +/- 1.96σ bands, we can conclude that the Sharpe ratio for the Minimum Variance 1% method is superior to the Market Cap method with statistical significance.

Showing all of the distributions graphically in one collection is visually difficult. The full set of results is shown in Table 5. The MC vs. EW distribution (orange curve) is also summarized in the top left entry in Table 5. The 'Diff' value of -0.012 is the mean difference in monthly Sharpe ratio (MC minus EW). The corresponding p-value of 0.38 means that we only have a 62% probability that the equal weight method has a superior Sharpe ratio than the market cap method. This is directionally intuitive but not statistically significant. Therefore, we cannot conclude that EW Sharpe ratio is above that of the MC Sharpe ratio.

For MV1 (blue curve), the 'Diff' value is -0.058 and the p-value is 0.007 which means that we have a 99.3% probability that the MV1 method has a superior Sharpe ratio than the MC method. This is directionally intuitive and statistically significant. We can therefore reject the null hypothesis.

2.8.2 Robustness Test - Results Summary

		EW	**IV**	**ER**	**MV1**	**MV5**	**MD1**	**MD5**
MC	Diff	-0.012	**-0.031**	**-0.031**	**-0.058**	**-0.056**	-0.031	-0.017
	p-value	0.380	0.026	0.026	0.007	0.018	0.091	0.404
EW	Diff		**-0.019**	**-0.019**	**-0.046**	**-0.044**	-0.019	-0.005
	p-value		0.000	0.001	0.022	0.046	0.188	0.760
IV	Diff			-0.000	-0.027	-0.025	-0.000	0.014
	p-value			0.967	0.112	0.198	0.998	0.415
ER	Diff				-0.027	-0.025	0.000	0.014
	p-value				0.090	0.180	0.992	0.386
MV1	Diff					0.002	**0.027**	**0.041**
	p-value					0.777	0.032	0.015
MV5	Diff						0.025	**0.039**
	p-value						0.091	0.021
MD1	Diff							0.014
	p-value							0.133

Table 5: Significance Testing of Monthly Sharpe Ratio Differences

In the table above, the **'Diff'** is the mean observed monthly difference of a bootstrap test between two Sharpe ratios. The number is constructed as the row header minus the column header. For the entry in the upper left of the table, this would be the observed mean Sharpe ratio of the MC method less the observed mean of the EW method. The number is negative as EW has a higher observed Sharpe ratio than the MC method. **p-value** is the significance level of the observed statistic. For a 95% confidence score the p-value must be less than or equal to 0.050; statistically significant difference values are noted with red or blue text. Values with 95% confidence are bold text while values with 90% confidence are non-bold text.

Using the method described in (Ledoit and Wolf 2008), for the differences of the 28

Sharpe ratio pairs, 10 of them are statistically significant at the 95% confidence level and 14 are significant at the 90% confidence level. In the simple ex-post point estimation of each Sharpe ratio over the 50 year period show in Table 17 and in Figure 3, all of the alternative allocation methods outperform the market cap allocation method. However, with the application of the robustness test, the same conclusions cannot be made. Observations for each method are reviewed below.

Market Cap (MC): While the MC method Sharpe ratio is dominated by all other methods in the simple ex-post comparison, using the advanced bootstrap method, the picture changes slightly. There is no statistical difference in Sharpe ratios between MC and either the EW or the MD approaches. At the 90% confidence level, MD1 Sharpe ratio is above that of MC.

Equal Weight (EW): In the hierarchy displayed in Figure 3, the EW method is dominated by all other methods except for MC. When viewing this through our bootstrap testing approach, most of these differences are confirmed with the exception of both MD1 and MD5 where a statistically significant claim cannot be made. Additionally, the MV5 Sharpe ratio is superior, but only at the 90% confidence level.

Inverse Volatility (IV): The table has already summarized the statistically significant claims of IV Sharpe being superior to MC and EW. However, no further claims can be made as the p-values all fail to indicate any true difference between IV and the remainder of the allocation methods.

Equal Risk (ER): Like the IV method, ER dominance is observed over MC and EW Sharpe ratios. At the 95% confidence level there can be no further claims about the differences with the remainder of the allocation methods. At the 90% confidence level however, MV1 dominates ER.

Minimum Variance (MV): MV1 and MV5 both dominate the MC and EW Sharpe ratios at the 90% level. There is no statistically significant difference between MV1 and MV5. In addition, both MV1 and MV5 dominate the MD1 and MD5 methods at the 90% confidence level.

Maximum Diversification (MD): Only the MC method appears to be dominated by the MD1 method and only at the 90% confidence level. Due to the wide standard error of the MD Sharpe ratios in the bootstrap tests, there are no other statistically significant claims that can be made about their dominance. Visual evidence of this can be seen in the width of the pink distributions in Figure 4.

Minimum Variance 1% (MV1) approach is the most frequent winner, dominating five of the seven other methods. Four of the superior Sharpe ratio claims can be made at the 95% confidence level while one can only be made at the 90% confidence level. Market cap is the clear laggard, reliably under-performing 4 of the 7 other methods at the 95% confidence level and an additional one at the 90% confidence level.

2.9 Conclusion

Over the past 50 years, many new methods have been developed for both security selection and capital allocation. With enormous amounts of capital seeking higher risk adjusted returns, work toward improving each of these stages is potentially valuable. Each stage in the formation of a portfolio can either add or subtract value.

Previous literature has established the capital allocation methods that were chosen for the panel in my study. Each of these research studies were noted during the description of the capital allocation methods earlier in this chapter. A subsequent study from (Chow et al. 2011) compared each of the strategies and showed that each of these strategies was almost entirely explained by the common factors of market, value, and size with a conclusion that implementation costs were a better way to examine their outcomes. A further study by (Clare, Motson, and Thomas 2013) showed that while all of the alternative capital allocation methods outperformed the market cap method, an investor should not be fooled by randomness. A fully random allocation method would also have performed quite well relative to the market cap method. Finally, a study by (Thomas, Clare, and Motson 2013) focused on using fundamental factors to create portfolios and these were not easily explained away by randomness.

This chapter extends the existing literature in a couple of ways. First the time frame of the previous studies ends in 2011 while my research extends until 2017, marking a full 50 years. Second, I apply a robustness check to the allocation methods to check whether the results observed in the point estimates of the Sharpe ratios are robust.

This chapter shows that there are capital allocation methods that can be universally employed (with a little expertise and computing power) which have provided a head start in the pursuit of improving risk adjusted returns. However, with the benefit of a more robust testing method, the clarity fades around a view that using just any systematic allocation method will help to improve the risk adjusted returns over a market capitalization method. All of the methods appear valid using a single sequence of historical returns but only a few stand up to more robust tests.

Minimum variance is nearly the exception as it statistically dominates four methods (MC, EQ, MD1, MD5) at the 95% confidence level and dominates ER at the 90% level. While not statistically significant, an 89% probability remains that MV1 is superior to the IV method. This is remarkable but confirms the point made earlier about Figure 2. The minimum variance allocation method uses no return data but fully derives the diversification benefits from the low and stable structure of the covariance matrix.

Even though we now see that 'all that glitters is not gold' for our capital allocation methods, this is just the first step in the analysis. So far, we have left security selection fixed, while examining the impact of capital allocation. In the next chapter, security selection is introduced and the relationship between the two stages begins to emerge and the conclusions shift further.

Chapter 3

Security Selection &

The Contours of Sharpe Ratios

3 Security Selection & The Contours of Sharpe Ratios

There are two stages in the Markowitz framework: security selection and capital allocation. In the prior chapter, we held security selection fixed (largest 1,000 stocks by market capitalization) while examining the hierarchy of Sharpe ratios across a panel of capital allocation methods. To understand the true nature of the hierarchy of Sharpe ratios, the analysis required a bit of complication with the introduction of the robustness tests. Now security selection is introduced to the process. By contrast, the effect of positive and negative security selection on Sharpe ratios is easy to intuit. By consistently picking better or worse stocks, we expect that all portfolio Sharpe ratios should rise or fall.

What becomes significantly more complex however, is understanding the interaction effect between any two processes. To examine the interaction between the two stages of portfolio formation, security selection is now applied across our panel of capital allocation methods. In the prior chapter, we made statistically robust conclusions about the risk adjusted performance differences between allocation methods. What we will now see in this step is a landscape of conclusions that shift as positive and negative security selection are applied.

The analysis begins by introducing a simple method for security selection. This simple approach is applied to a single capital allocation method (market capitalization) to isolate the interaction effect of a single variable against the panel of fixed allocation methods. Robustness checks are applied. Then the same security selection method is applied to all capital allocation methods at the same time. This has the effect of creating a contour of Sharpe ratio differences across the landscape of positive and negative security selection tilt. Robustness checks are then applied across each contour and we observe that initial conclusions can vary significantly.

The key finding from this chapter is that the optimal choice of a capital allocation method depends upon the ability to select stocks. Poor security selection skills can be mitigated by protective allocation methods while positive security selection skills can be amplified by using the right allocation method. The new conclusions are verified using the same robust bootstrap method to verify differences between Sharpe ratios and their potential improvements.

3.1 Security Selection

Aside from fully discretionary (and thereby non-replicable) security selection methods, there are a myriad of systematic security selection techniques. As the goal of this analysis is to study the interaction effect, I do not pursue a path to prove which of these methods is better than any other. Instead, I introduce a simple method which, at the limit, will have the most pronounced effects (positive and negative) on the selection stage of portfolio formation.

Security selection is simply a process to move from the full 1,000 stocks available each year in the data set to a smaller set. There are likely infinite ways that to create the security selection process in order to observe this interaction. Fundamental characteristics of stocks could be used to tilt the security set upward or downward (size, value, momentum, etc.) with a nearly unlimited combination of parameters. Purely quantitative approaches could also be employed for this purpose (momentum, trend filters, etc.) Fundamental and quantitative approaches could be combined and an unlimited number of parameters for each method could be used.

Following any of these methods may systematically tilt the security set upward or downward which would be helpful in observing the effect on Sharpe ratios. However, following any of these methods or combinations of methods would create a bias in why a certain set of factors was selected with the necessary accompanying parameters. For this reason, security selection with near perfect foresight was chosen as a boundary condition that requires no parameters and no intended biases.

A simplistic, yet unrealistic method for selecting stocks is still appropriate here as the purpose of the analysis is not to suggest a way to pick stocks, but rather to study the interaction effect between the selection and allocation stages of portfolio formation.

3.1.1 Security Selection with Near Perfect Foresight

Markowitz famously wrote about the process of selecting a portfolio:

> *"The first stage starts with observations and experience and ends with beliefs about the future performances of available securities." (Markowitz 1952)*

From this vantage point, perfect foresight would allow us to choose the single best performing security for the next period and put all of our money into that one holding. But he also cautioned against holding a single security.

> *"The hypothesis (or maxim) that the investor does (or should) maximize discounted return must be rejected. Diversification is both observed and sensible; a rule of behavior which does not imply the superiority of diversification must be rejected both as a hypothesis and a maxim." (Markowitz 1952)*

If the objective of portfolio formation were simply to achieve the highest returns possible, irrespective of risk, then perfect foresight would select the one stock with the highest returns. However, Markowitz clearly rules this out as a practical path and sets in place the superiority of diversification.

Diversification is only possible with more than one security. This analysis seeks to understand how the Sharpe ratios from a panel of allocation methods change in the presence of security selection (positive and negative). So to form portfolios from sets of securities with better and worse performance, we need a method to pick stocks each year to feed into the allocation methods and observe the resulting Sharpe ratios.

As the existing literature points out, there is a vast array of methods espoused for selecting stocks. These can be from any combination of discretionary or systematic, and from fundamental and quantitative. Each of these methods can select from a wide variety of inputs (financial statement ratios, relative momentum, etc.) and parameters (1 year look-back, rolling 3 years, etc.). Any method that we might use in our analysis can be adapted with a slight change in parameter set or inclusion of an additional data point. And any method that we use may work well in a particular year and poorly in a later year.

For this reason, I introduce the concept of near-perfect foresight. It is not achievable in practice, but it is simple to implement, reliably creates the stock return bias that is needed for answering my particular questions, and leaves intact the requirement of diversification that Markowitz established. Near-perfect foresight is a boundary condition method in that it is the best outcome that can be achieved while preserving diversification. No other method of security selection could outperform near-perfect foresight and retain the same potential for diversification.

So if perfect foresight implies picking one single best winner, then my versions of near-perfect foresight would be identifying in advance a range of stocks with the biggest gains or losses in the following year and removing them from the set of available securities. This method leaves intact the potential for diversification by continuing to work with a large set of available securities while reliably improving or degrading the return profile of the potential portfolios.

The approach is implemented by using a one year look ahead when forming the portfolios. As before, it begins with the 1,000 largest capitalization stocks available in the data set at the end of the year. I then look ahead a full year and rank the performance of each stock that is available for the portfolios. Portfolios are formed with advance knowledge by systematically removing the best or worst performing stocks based on performance in the following year.

3.2 Positive Selection & Market Capitalization

To keep the first step of this analysis simple, I begin by applying only positive security selection to a single capital allocation method, market capitalization. Using the near perfect foresight approach, I look ahead one year and remove the worst performing stocks. This is done in batches of 5 stocks up to a maximum of 100 stocks being removed. Each year when the portfolios are formed, the one year look ahead and screening is implemented to create a positive tilt in returns and Sharpe ratios.

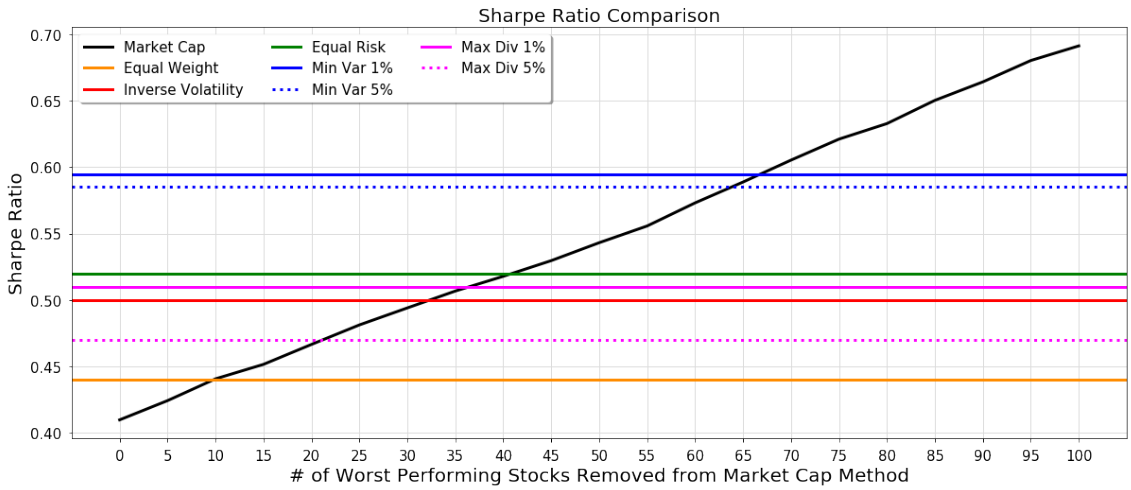

Figure 5: Upward tilted market capitalization weighting method

The figure above documents the increase in MC Sharpe ratio when removing the worst performing stocks using a look-ahead tilt. Stocks are removed in batches of 5 up to 100. The 5 worst performing are removed first, then the worst 10, etc.

The positive security selection tilt is only applied at this point to the MC method while all other allocation methods remain formed with the fixed set of 1,000 largest capitalization stocks. Figure 5 shows the impact of this process on the market cap line (black) while the other allocation lines remain static. This result is as expected.

As poor performing stocks are eliminated from the market cap weighted portfolio, Figure 5 shows the rise in the Sharpe ratio relative to the other allocation methods which do not have any security selection applied. This increase in Sharpe ratios is not precisely linear but rises as expected. However, as demonstrated in the previous chapter, this comparison of the point estimates of Sharpe ratios is too simplistic and lacks the robustness of a studentized bootstrap method of comparison.

For example, Figure 5 shows that the MC Sharpe ratio is above the EW Sharpe ratio once the bottom 15 stocks are eliminated each year. However, from the robustness checks earlier in Table 5, it can't be claimed that EW method was better than MC method using the full 1,000 stock set. In order to determine a statistically significant improvement in Sharpe ratios, both a simplistic security selection process and the studentized bootstrap method of comparing differences in Sharpe ratios must be applied.

Table 6 combines security selection with a robust Sharpe ratio test. The columns of the table represent the panel of alternative allocation methods while the rows relate show the MC method with increasing levels of perfect foresight security selection applied. The identification of MC+10 means that the returns of the MC method have been tilted higher ('+') by removing the worst 10 performing stocks each year. 'Diff' values are constructed as the row method minus the column method. For example the first entry in the upper left of -0.012 implies that the EW method Sharpe ratio is higher than the MC method Sharpe ratio. Positive values which are statistically significant are denoted with blue text and negative values which are statistically significant are denoted with red text. Values which are significant at the 95% level are also placed in bold text while

		EW	IV	ER	MV1	MV5	MD1	MD5
MC	Diff	-0.012	**-0.031**	**-0.031**	**-0.058**	**-0.056**	-0.031	-0.017
	p-value	0.380	0.026	0.022	0.007	0.018	0.091	0.404
MC+10	Diff	-0.003	-0.022	-0.022	**-0.049**	**-0.047**	-0.022	-0.008
	p-value	0.835	0.114	0.105	0.017	0.036	0.229	0.697
MC+20	Diff	0.005	-0.015	-0.015	**-0.041**	-0.039	-0.014	-0.000
	p-value	0.720	0.294	0.285	0.048	0.078	0.420	0.978
MC+30	Diff	0.013	-0.007	-0.007	-0.033	-0.031	-0.007	0.007
	p-value	0.337	0.633	0.636	0.107	0.177	0.709	0.723
MC+40	Diff	0.020	0.000	-0.000	-0.027	-0.024	0.000	0.014
	p-value	0.141	0.984	0.989	0.201	0.275	0.989	0.484
MC+50	Diff	**0.027**	0.008	0.008	-0.019	-0.017	0.008	0.022
	p-value	0.048	0.574	0.570	0.345	0.448	0.675	0.299
MC+60	Diff	**0.036**	0.016	0.016	-0.010	-0.008	0.017	0.030
	p-value	0.008	0.217	0.202	0.614	0.722	0.358	0.141
MC+70	Diff	**0.045**	**0.026**	**0.026**	-0.001	0.001	0.026	0.040
	p-value	0.002	0.044	0.050	0.964	0.960	0.149	0.053
MC+80	Diff	**0.053**	**0.034**	**0.034**	0.007	0.009	0.034	**0.048**
	p-value	0.000	0.010	0.009	0.730	0.679	0.057	0.023
MC+90	Diff	**0.062**	**0.043**	**0.043**	0.016	0.018	**0.043**	**0.057**
	p-value	0.000	0.002	0.002	0.435	0.413	0.018	0.006
MC+100	Diff	**0.070**	**0.051**	**0.051**	0.024	0.026	**0.051**	**0.065**
	p-value	0.000	0.000	0.000	0.227	0.225	0.005	0.003

Table 6: Monthly Sharpe Ratio Comparisons With Security Selection: Market Cap method

In the table above, the **'Diff'** is the mean observed monthly difference of a bootstrap test between two Sharpe ratios. The number is constructed as the row header minus the column header. For the entry in the upper left of the table, this would be the observed mean Sharpe ratio of the MC method less the observed mean of the EW method. The number is negative as EW has a higher observed Sharpe ratio than the MC method. **p-value** is the significance level of the observed statistic. For a 95% confidence score the p-value must be less than 0.05. Table cells where the difference values are positive and statistically significant are colored in blue. Table cells where the difference is negative and statistically significant are colored in red. Bold cells in color indicate 95% confidence while non-bold cells in color indicate 90% confidence. Each row compares the allocation method with a certain number of poorly performing stocks removed (e.g. MC+10 is the market cap method tilted upward by removing the 10 worst performing stocks from the stock universe each year.)

values significant only at the 90% level are non-bold.

The application of a robustness check makes evident that for the MC method to produce a higher Sharpe ratio than EW method and be statistically significant, the original observation of removing the bottom 15 stocks each year is insufficient. Instead it would be necessary to remove the worst performing 50 stocks each year. The feasibility of achieving this level of near perfect foresight in security selection with 50 years of continuous persistence considerably undermines the notion that security selection can reliably be implemented on a market cap weighted portfolio in pursuit of producing superior risk adjusted returns.

Observations about the panel of alternative allocation methods are made below:

- **Equal Weight (EW)**: Without security selection, a simple ex-post point estimate of the Sharpe ratio of the EW approach appears superior to the MC method. Yet no such claim can be made under the approach of a robust test of their difference. Even in the presence of relatively high levels of perfect security selection, EW is statistically indistinguishable from the MC method. Only after the application of removing the 50 worst performing securities from MC every year for 50 years in a row do we observe Sharpe ratio improvements of MC that are reliably better than EW without security selection.

- **Inverse Volatility (IV)**: The IV method has shown a statistical robustness over the MC method without any security selection. In the presence of lower levels of security selection, the two methods become indistinguishable until at least 70 or 7% of securities are correctly removed every year for 50 years in a row. Only at this heightened level of prescience does the MC method significantly outperform the IV without security selection.

- **Equal Risk (ER)**: Identical to the IV method, the ER method has shown a statistical robustness over the MC method without any security selection. In the presence of lower levels of security selection, the two methods become indistinguishable until at least 70 or 7% of securities are correctly removed from MC every year for 50 years in a row.

- **Minimum Variance (MV)**: This allocation method is the most robust against an MC portfolio at any observed level of security selection. In addition to statistically dominating the MC method without security selection, the MV methods retain dominance all the way through the perfect elimination of the 20 worst securities each year for the MC method. Remarkably, even at very high levels of security selection (from MC+20 up to MC+100), there is no statistically significant difference in the Sharpe ratios. The MV allocation method Sharpe ratio can not be dominated by the MC method even at the highest level of security selection used in our study.

- **Maximum Diversification (MD)**: At 90% confidence, the MD1 method appears superior to the MC method without any security selection expertise. The two methods are indistinguishable starting from lower levels of security selection (MC+10) up through relatively high levels (MC+70) of security selection. The wide range over which the two methods are indistinguishable from each other is mainly due to the high standard error that the MD Sharpe ratios display in bootstrap testing. Only at the higher levels of security selection (MC+70 and higher) do the MC methods with positive security selection begin to dominate the Sharpe ratio of MD methods without security selection.

3.3 Simultaneous Security Selection tilt

If security selection skill is available to the portfolio manager, it can be applied in combination with any capital allocation method. In the previous section, the impact of security selection tilt on the MC allocation method appeared to be relatively linear. However, in the next step I apply the same 'near perfect foresight' method to all allocation methods simultaneously. By extending the security selection tilt across the full panel of allocation methods, a non-linear impact on Sharpe ratios is documented. Simultaneously applying security selection and capital allocation methods creates a landscape of Sharpe ratio differences which highlight the interaction effect created by the combination of selection and allocation.

3.3.1 Simultaneous Positive Security Selection

As before, the process begins with the full set of 1,000 largest market capitalization stocks each year. Using a one year look ahead on returns, stocks are removed from the full set in batches of 5. To keep moving ahead in small steps, only positive security selection tilt is applied as before, but now it is applied across the full panel of capital allocation methods.

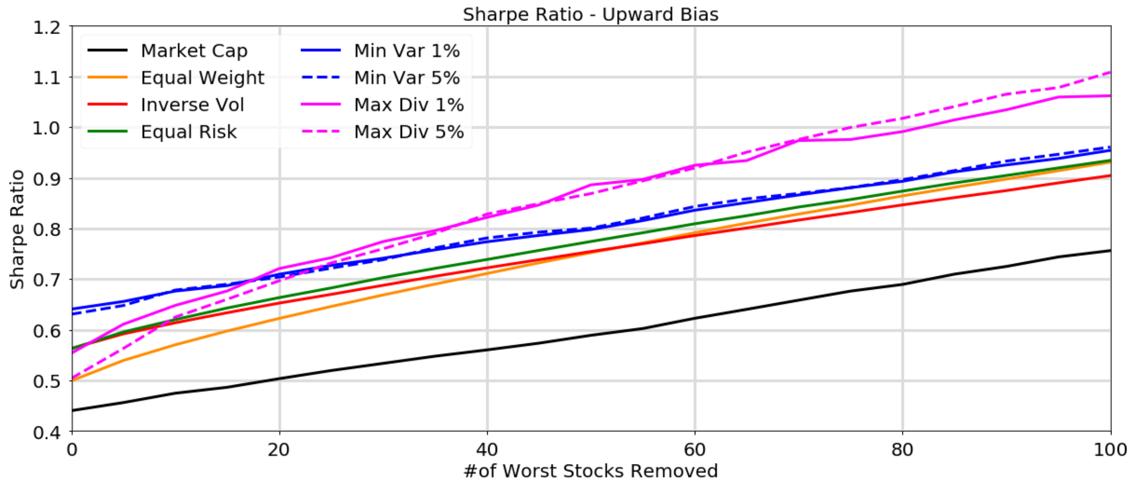

Figure 6: Sharpe ratios with Positive Security Selection tilt
The figure above documents the effect that a positive security selection tilt has on the Sharpe ratios across the full panel of allocation methods. Upward security selection tilt is created by removing the worst performing stocks each year. Statistics are reflected across the period of study, 1968 - 2017.

Figure 6 highlights the non-linear impact on Sharpe ratios despite the application of identical security selection tilt. For example, the MD methods (pink lines) sit near the bottom of the hierarchy in the absence of security selection (far left of graphic). However, as the positive security selection tilt increases (moving toward the right), the MD methods exhibit amplified Sharpe ratios and begin to outperform the Sharpe ratios for all other methods.

Several additional observations appear visually. Without security selection, equal weight appeared to outperform market cap. This claim was rejected when robustness checks were applied. Now, in the presence of positive security selection, the spread between these two lines (MC - black, EW - Orange) begin to diverge and it might be possible to make new claims about their relationship. This will be examined with robustness checks across this new landscape.

At each point along the landscape of security selection tilt, the original claims of robust differences between Sharpe ratios must be reexamined for each pair. But before robustness checks are applied, let's extend this approach to examine the landscape across a negative security selection tilt.

3.3.2 Simultaneous Negative Security Selection

Negative security selection is crafted using the same method as for positive security selection tilt in the previous section. Instead of removing batches of 5 worst performing stocks each year, I remove the best performing stocks in batches of 5 up to a total of 100. Another difference to note in Figure 6 is that the horizontal axis started at zero (left side) and moved upward toward the right. To reflect the opposite action in this section, I reverse the axis to demonstrate the negative impacts of poor security selection.

As with positive security selection tilt, there are some interesting observations which

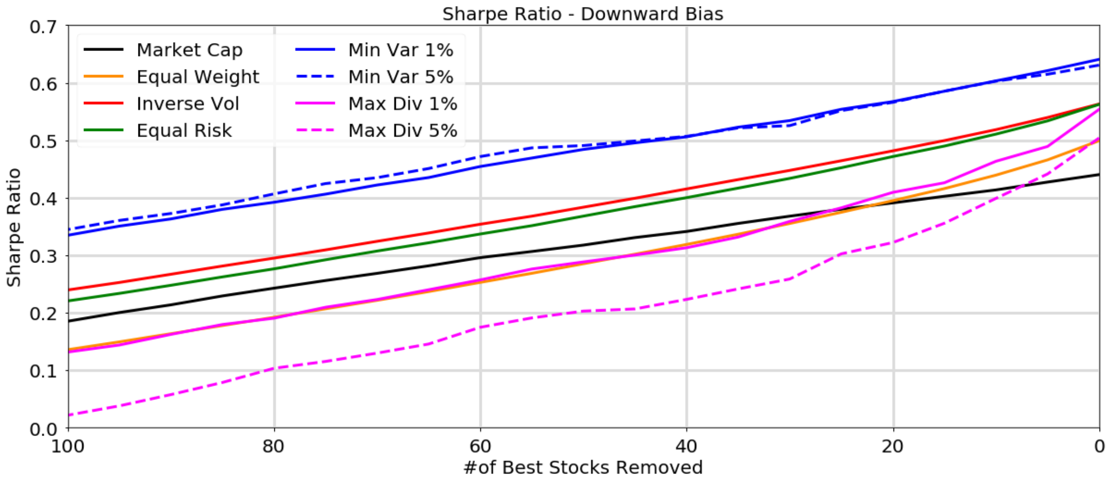

Figure 7: Sharpe ratios with Negative Security Selection tilt
The figure above documents the effect that a negative security selection tilt has on the Sharpe ratios across the full panel of allocation methods. Downward security selection tilt is created by removing the best performing stocks each year. Statistics are reflected across the period of study, 1968 - 2017.

appear visually, and these also will be checked for robustness in a following section.

Maximum diversification methods (pink lines) exhibit accelerating under-performance in the presence of negative security selection skills. They appear to markedly underperform all other methods relatively early in the landscape. Equal risk (green) and inverse volatility (red) methods have an early advantage on market cap method, but as negative security selection increases, this lead is diminished. Minimum variance (blue) methods appear to maintain their lack of sensitivity to poor security selection.

3.3.3 Security Selection - Full Landscape

Figure 8 combines both positive and negative security selection into a single visual. At the center of the graph (zero on the horizontal axis), the original hierarchy of Sharpe ratios is displayed. Moving to the right shows the shifting hierarchy in the presence of positive security selection while moving to the left shows the shifting hierarchy in the presence of negative security selection.

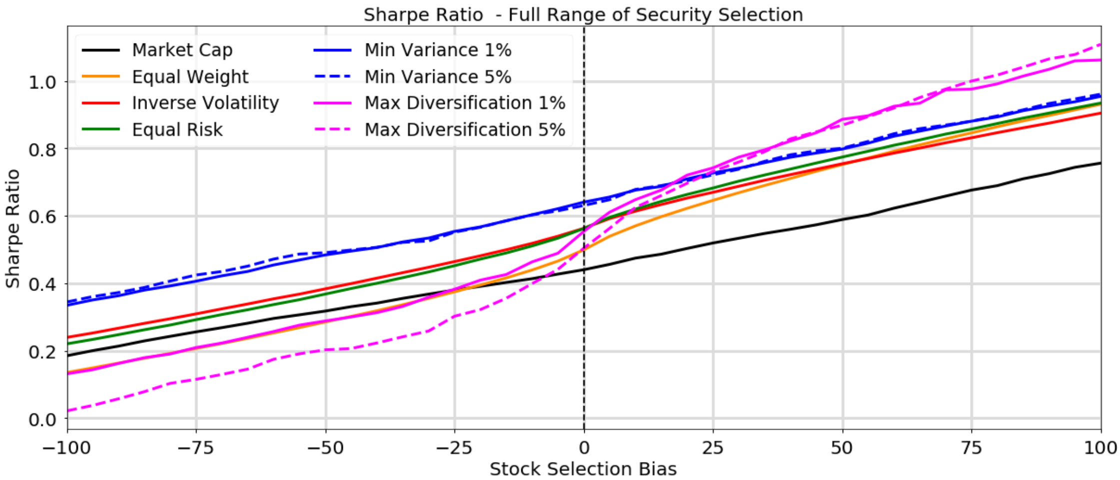

Figure 8: Sharpe Ratios with Positive and Negative Security Selection Tilt
The figure above documents the effect of both positive and negative security selection on the Sharpe ratios across the full panel of allocation methods. Statistics are reflected across the period of study, 1968 - 2017.

There are two types of shifts that we observe in the panel of allocation method Sharpe ratios. One type of shift has a non-linear element, while the other type of shift is a difference in the slope of the nearly linear changes. Each of these acts to change the hierarchy at the endpoints of the security selection tilt range.

3.4 Individual Contours - Examples

This chapter highlights the instability of the hierarchy of Sharpe ratios in the presence of security selection. To show this effect clearly, robustness checks must be applied to each pair of Sharpe ratios along the entire landscape of security selection tilt. In the following subsections, I highlight a few examples of the shifting contours in the pairs of Sharpe ratios. A full set of visual contours is provided in Appendix 2.

3.4.1 Market Cap vs. Equal Weight

Let's start with a first example. Back in the previous chapter, Figure 4, the top left chart showed the robustness check of the EW Sharpe ratio vs. the MC Sharpe ratio. While EW appeared to have a Sharpe ratio superior to the MC allocation method, due to the width of the confidence band of a two tailed bootstrap test, we could not reject the null hypothesis that the Sharpe ratios are equal.

Applying a range of security selection across both the market cap weighted portfolio and the equal weighted portfolio has a significant impact on the probability that the equal weight method is superior to the market cap method.

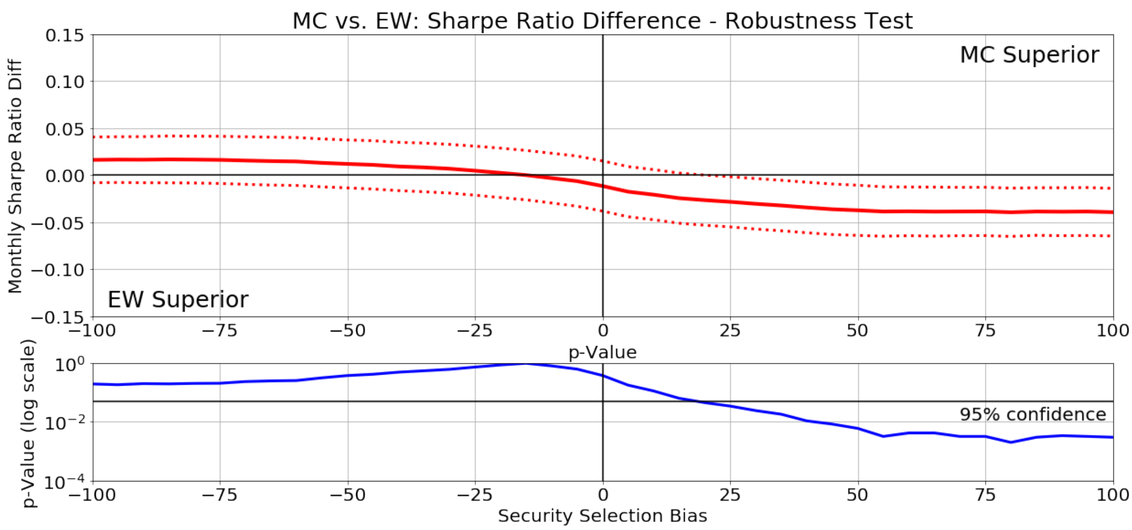

Figure 9: Sharpe ratio contour: MC vs. EW
The top panel above documents the probability distribution of the difference in monthly Sharpe ratios between two allocation methods across a range of security selection tilt. The solid line is the estimated value while the dashed lines represent a 95% confidence interval. The vertical axis is the monthly Sharpe ratio difference of the MC method minus the monthly Sharpe ratio of the EW method. When the zero horizontal line is contained within the 95% probability distribution, then the null hypothesis (the true difference is zero) can not be rejected. The bottom panel gives the p-value for the two tailed test. When the p-value is below 0.05, there is at least a 95% confidence that the true value of the difference in Sharpe ratios is not zero.

The tests previously conducted on the difference in monthly Sharpe ratios between the market cap and equal weight allocation methods could not reject the null hypothesis and had to conclude that the true difference in monthly Sharpe ratios was not statistically different than zero. However, now when the probability distribution is estimated across the landscape of security selection tilts, a different conclusion emerges.

By looking at the full range of security selection tilts, we see that the initial inability to conclude a difference in Sharpe ratios gives way to a significant conclusion in the presence of sufficient security selection skill. It is now possible to conclude that once a manager is able to accurately remove 20 or more of the worst performing stocks each year from the basket of 1,000 large cap stocks, the Sharpe ratio of the equal weight allocation method rises more quickly than the market cap method. The improvement in Sharpe ratios is non-linear and beyond the 20 stock selection level, the difference in Sharpe ratios is now statistically significant.

With this particular allocation method pair, no significant conclusions can be made in the presence of negative security selection skills. While the estimated difference increases, indicating that market cap is becoming superior to equal weight, the zero horizontal line stays within the 95% confidence interval which means that the null hypothesis cannot be rejected.

This is just a first example of the non-invariance displayed by the hierarchy of Sharpe ratios when security selection is applied. In simple terms it means that in the absence of security selection, there is no statistically significant improvement in a portfolios outcome from choosing the EW allocation method rather than the MC allocation method. However, if a manager can apply a certain level of positive security selection, the choice of the equal weight allocation method can improve the outcome of the Sharpe ratio with a high degree of statistical significance.

3.4.2 Market Cap vs. Minimum Variance 5%

In the previous section we saw an example of a strategy pair without a significant conclusion at the beginning. But by applying positive security selection, the conclusion changed and we saw the superiority of the equal weight allocation method. This type of shift in conclusion is not always the same. Each pair of allocation methods reacts differently to the application of security selection tilt.

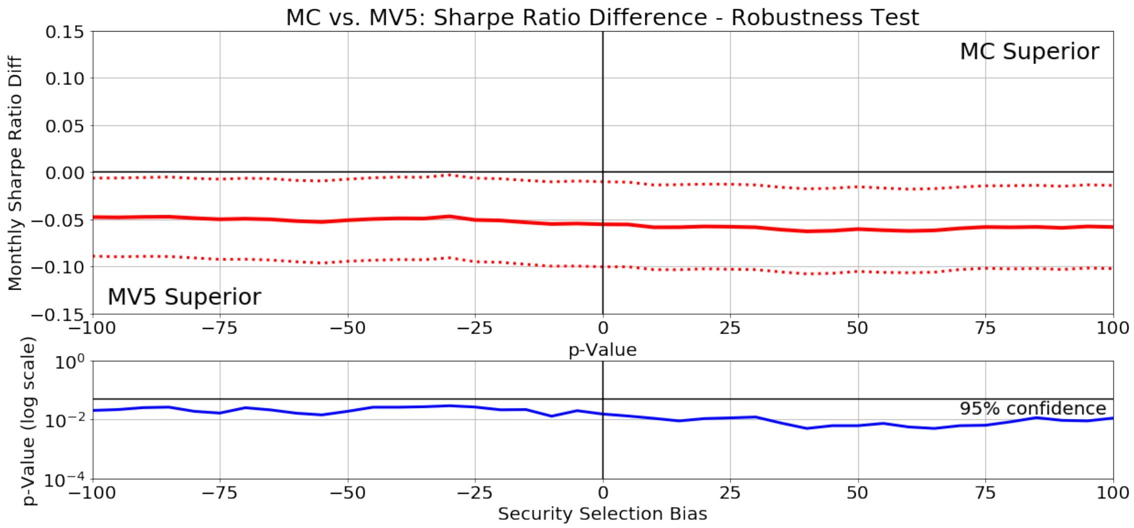

Figure 10: Sharpe ratio contour: MC vs. MV5%
The top panel above documents the probability distribution of the difference in monthly Sharpe ratios between two allocation methods across a range of security selection tilt. The solid line is the estimated value while the dashed lines represent a 95% confidence interval. The vertical axis is the monthly Sharpe ratio difference of the MC method minus the monthly Sharpe ratio of the MV5% method. When the zero line is contained within the 95% probability distribution, then the null hypothesis (the true difference is zero) can not be rejected. The bottom panel gives the p-value for the two tailed test. When the p-value is below 0.05, there is at least a 95% confidence that the true value of the difference in Sharpe ratios is not zero.

In Figure 10, the allocation pair shows a different contour across the landscape of security selection. In the absence of security selection, we can make a statistically significant conclusion that the MV5 method is superior to the market cap method. And with both positive and negative security selection applied, this conclusion does not change. In a head to head match up across the full range of security selection, our initial conclusion is invariant. The MV5 method wins over across the entire range of security selection tilt.

3.4.3 Market Cap vs. Maximum Diversification 5%

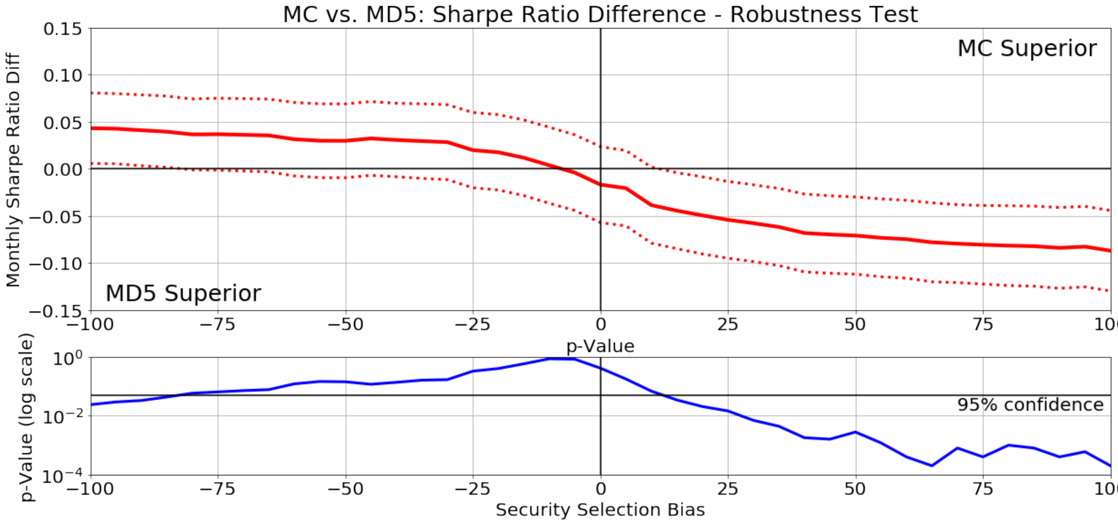

Figure 11: Sharpe ratio contour: MC vs. MD5%
The top panel above documents the probability distribution of the difference in monthly Sharpe ratios between two allocation methods across a range of security selection tilt. The solid line is the estimated value while the dashed lines represent a 95% confidence interval. The vertical axis is the monthly Sharpe ratio difference of the MC method minus the monthly Sharpe ratio of the MD5% method. When the zero line is contained within the 95% probability distribution, then the null hypothesis (the true difference is zero) can not be rejected. The bottom panel gives the p-value for the two tailed test. When the p-value is below 0.05, there is at least a 95% confidence that the true value of the difference in Sharpe ratios is not zero.

The final example shown is one in which the conclusion inverts with significance at the ends of our selection tilt range. Figure 11 starts with no significant conclusion in the absence of security selection. With poor security selection however, the market cap allocation method become significantly superior at the extremes of the range examined. This conclusion shifts from neutral to statistically significant around -90.

In the presence of positive security selection, the conclusion inverts entirely. At +15 in security selection tilt, the maximum diversification method begins to be significantly superior to the market cap allocation method.

The maximum diversification allocation methods amplify the security selection tilt. It is interesting that positive security selection appears to benefit the MD5 methods more than negative security selection impairs the Sharpe ratio difference of this match-up vs. MC. This particular pair of allocation methods is one of the very few that is observed with complete inversions in conclusion in the presence of security selection.

3.5 Landscape Summary

In the preceding sections, a few examples of the landscape contours have been shown which demonstrate the types of shift in conclusion that occur in the presence of security selection tilts. The full set of 28 visuals for each pair of allocation methods that are tested follow at the end of the chapter.

One way to summarize all of the shifting conclusions is provided in the combination of Table 7 and Table 8. Table 7 displays changes in conclusion in the presence of positive security selection tilt while Table 8 displays changes in conclusion in the presence of negative security tilts.

A complete explanation of the entries in the tables is required. The first entry in the upper left of Table 7 compares the Sharpe ratio difference created by the row (MC) minus the column (EW). The first row of each entry is the difference ('Diff') observed at the starting point of zero security selection tilt. In this case the entry of '?/-' indicates that there is no statistical significance ('?') observed at the starting point, but the conclusion shifts to negative ('-') with statistical significance. A negative conclusion means that the row (MC) minus the column (EW) is now a negative value and statistically significant at the 95% confidence level.

		EW	IV	ER	MV1	MV5	MD1	MD5
MC	Diff	?/-	-	-	-	-	?/-	?/-
	Level	+20					+5	+10
EW	Diff		-/?	-/?	-/?	-/?	?/-	?/-
	Level		+25	+45	+10	+5	+20	+60
IV	Diff			?/-	?	?	?/-	?/-
	Level			+45			+40	+65
ER	Diff				?	?	?/-	?/-
	Level						+40	+70
MV1	Diff					?	+/?/-	+/?
	Level						+5/+95	+5
MV5	Diff						+	+/?
	Level							+5
MD1	Diff							?
	Level							

Table 7: Changes in the significance test of monthly Sharpe ratio differences in the presence of positive security selection tilt

In the table above, the **'Diff'** is the robust conclusion for the difference of monthly Sharpe ratios for two allocation methods. The table is constructed so that the value is row minus column. Where there are two values (e.g. ?/-) this indicates the initial conclusion (?) followed by the ultimate difference (-) at some level of security selection tilt. The **'Level'** indicates at which level of security selection the 'Diff' switches. If there is no value in the 'Level' row, then the initial conclusion remains constant throughout the range of security selection.

		EW	IV	ER	MV1	MV5	MD1	MD5
MC	Diff	?	-/?	-/?	-	-	?	?/+
	Level		-10	-10				-85
EW	Diff		-	-	-	-	?	?
	Level							
IV	Diff			?/+	?/-	?/-/?	?/+	?/+
	Level			-95	-45	-55/-85	-40	-15
ER	Diff				?/-	?/-	?/-	?/-
	Level				-10	-40	-40	-10
MV1	Diff					?	+	+
	Level							
MV5	Diff						+	?/+
	Level							-5
MD1	Diff							?/+
	Level							-15

Table 8: Changes in the significance test of monthly Sharpe ratio differences in the presence of negative security selection tilt

In the table above, the **'Diff'** is the robust conclusion for the difference of monthly Sharpe ratios for two allocation methods. The table is constructed so that the value is row minus column. Where there are two values (e.g. ?/-) this indicates the initial conclusion (?) followed by the ultimate difference (-) at some level of security selection tilt. The **'Level'** indicates at which level of security selection the 'Diff' switches. If there is no value in the 'Level' row, then the initial conclusion remains constant throughout the range of security selection.

The second row of this cell entry shows the level of positive security selection tilt at which the conclusion shifts. In this case, EW becomes superior to MC allocation method and statistically significant once we achieve a +20 level of positive security selection tilt occurs.

Some of the conclusions never change. For example, the second entry in the top row examines MC minus IV. In this case, IV is superior to MC at all levels of positive security selection. Therefore, the entry is '-' and has no level at which the conclusion shifts. This is also the case for entries such as IV minus MV1 (3rd row, second entry). In this case, there is no statistically significant conclusion ('?') in the absence of security selection or anywhere along the range of positive security selection tilt.

A few entries in the tables have conclusions that shift more than once. The MV1 minus MD1 begins with a '+' which indicates that minimum variance is superior. It then shifts to inconclusive ('?') at +5 of security selection tilt. Finally it shifts to minimum variance being inferior ('-') at the +95 level of positive security selection tilt.

Of the 28 pairs examined, only 11 pairs showed consistency in the original conclusion based on no security selection tilt. The majority, 17 of 28 pairs, had a statistically significant shift in their shape across the landscape of positive security selection tilt.

Table 8 shows the same information with the application of negative security selection tilt. On the negative security selection tilt side of the landscape, half of our conclusions shift while half stay the same.

Only 3 pairs remain with their original conclusions throughout the full range of security selection tilt. MV1 and MV5 always outperform the MC method, while MV1 vs MV5 is consistently inconclusive.

3.6 Conclusion

This chapter creates a bridge between two large strands of existing literature. Following (Markowitz 1952) who laid out the two stages of portfolio formation, one strand of literature has followed the route to finding better ways to select securities while the other strand has followed the route to improve upon capital allocation method. This chapter spans these two paths in the literature by examining how security selection can unevenly impact various capital allocation methods.

This chapter began by creating a simple yet robust method of security selection as a tool for extracting the interaction effect between the stages of portfolio formation. The 'near-perfect foresight' method allowed us to highlight the non-invariance of the hierarchy of Sharpe ratios in the presence of security selection. A variety of contours was documented across the landscape of security selection tilt.

In the prior chapter we established an initial view of the hierarchy of Sharpe ratios, yet with the application of a robust test found that 'all that glitters is not gold'. Now we see a similar shift in our conclusions. When security selection tilt is applied, the contour of each pair reacts differently and can further shift our prior conclusions. This is tested in the same robust framework as before.

The key finding of this chapter is that the optimal choice of a capital allocation methods depends upon the ability to pick stocks. Poor security selection skills can be mitigated by protective allocation methods while positive security selection skills can be amplified by using the right allocation method.

We have now taken a few steps forward in understanding the drivers of optimal portfolio construction. In the next chapter, we use 'near-perfect foresight' to highlight the key drivers of optimal Sharpe ratio portfolios.

Chapter 3

Appendix

3.7 Sharpe Ratio Contours

3.7.1 Description of Contours

The figures that follow are a collection of Sharpe ratio contours across a range of security selection bias. These figures are referenced in and explained in Chapter 3. Each figure in this section follows the same formatting guidelines which provide for clarity.

The horizontal axis gives a range of security selection bias and is the same for both panels. The center point of the horizontal axis reflects no selection tilt while to the left is a negative selection tilt and to the right is a positive selection bias. Security selection tilt is based on the 'near perfect foresight' approach described in Chapter 3.

The top panel is a contour with the solid mid-line giving the point estimate of the monthly Sharpe ratio difference while the dashed line bands form the two tailed confidence interval at 95%.

The lower panel gives the p-value of the test and values below the solid black line indicate that the test has achieved at least a 95% confidence. The null hypothesis is that the true difference of Sharpe ratios between the pair of allocation methods is zero. p-Values below 0.05 allow us to reject the null hypothesis and conclude a difference exists.

A subsection in this appendix has been created for each row in the matrices shown in Table 7 and Table 8.

3.8 Market Cap vs. Panel of Allocation Methods

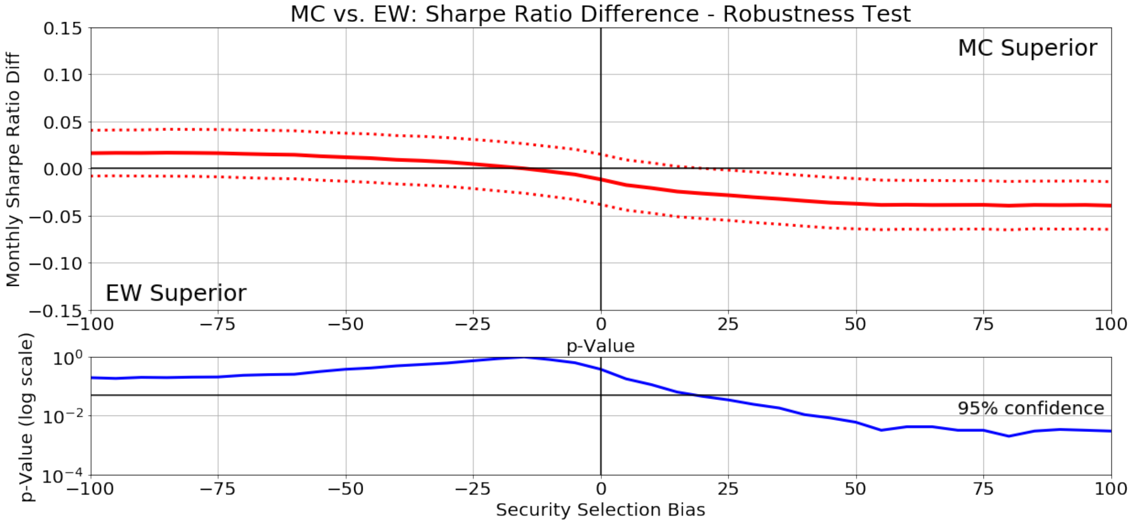

Figure 12: Sharpe ratio contour: MC vs. EW

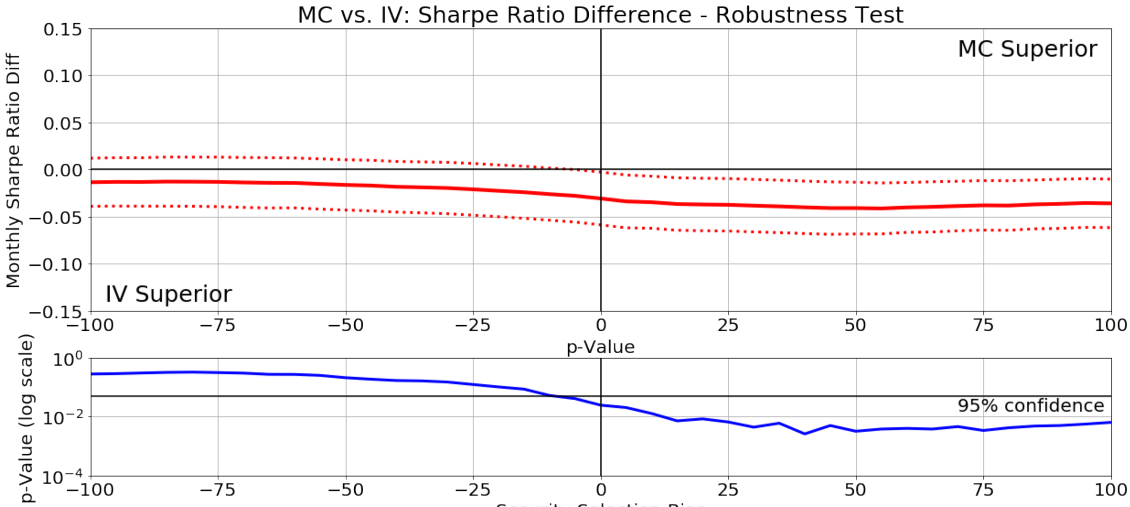

Figure 13: Sharpe ratio contour: MC vs. IV

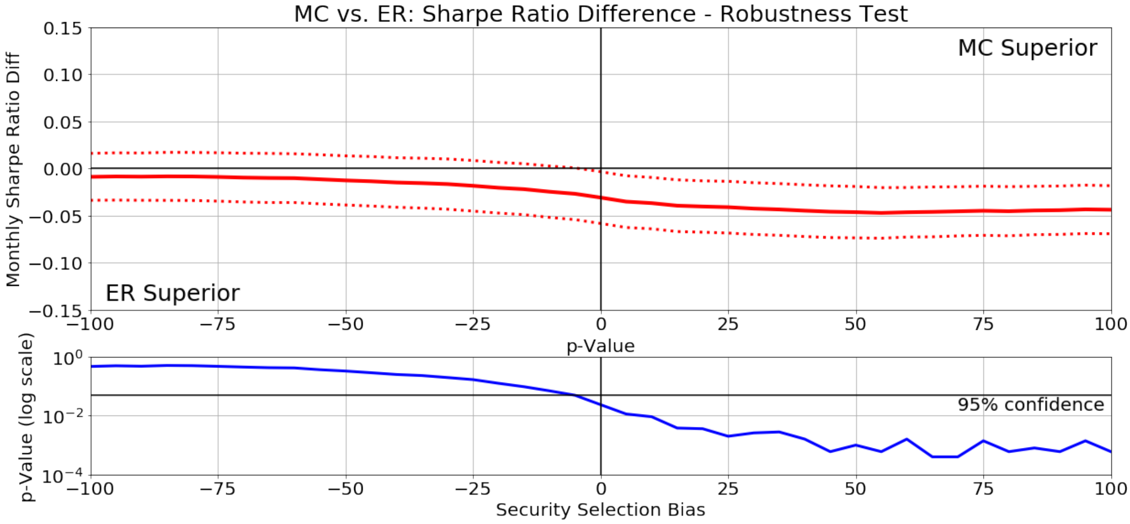

Figure 14: Sharpe ratio contour: MC vs. ER

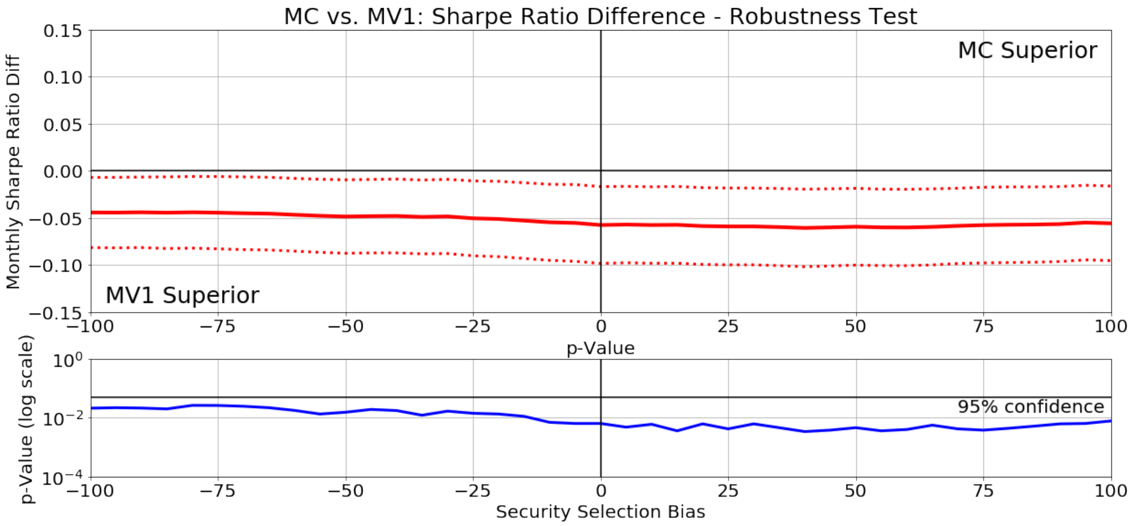

Figure 15: Sharpe ratio contour: MC vs. MV1

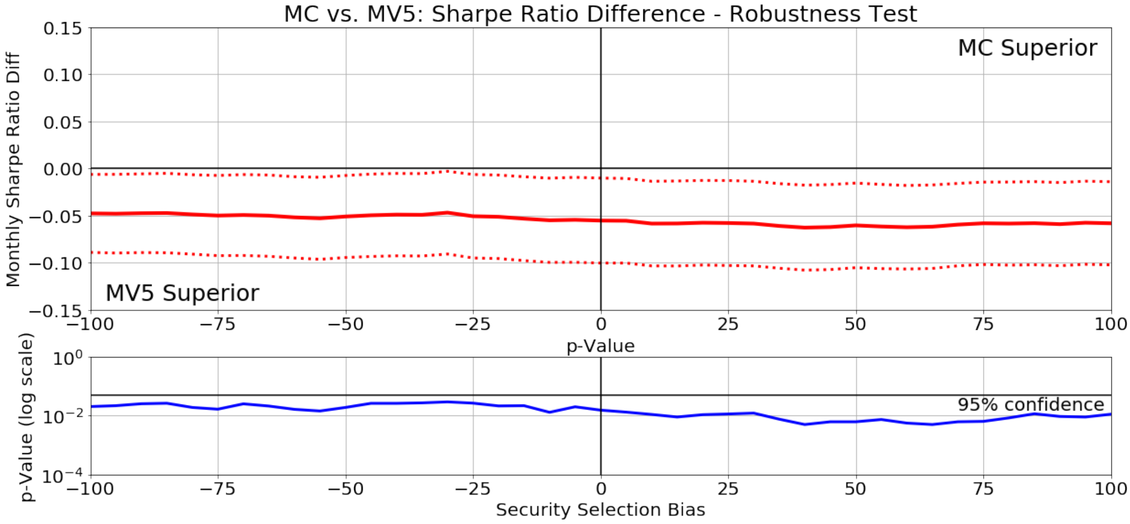

Figure 16: Sharpe ratio contour: MC vs. MV5

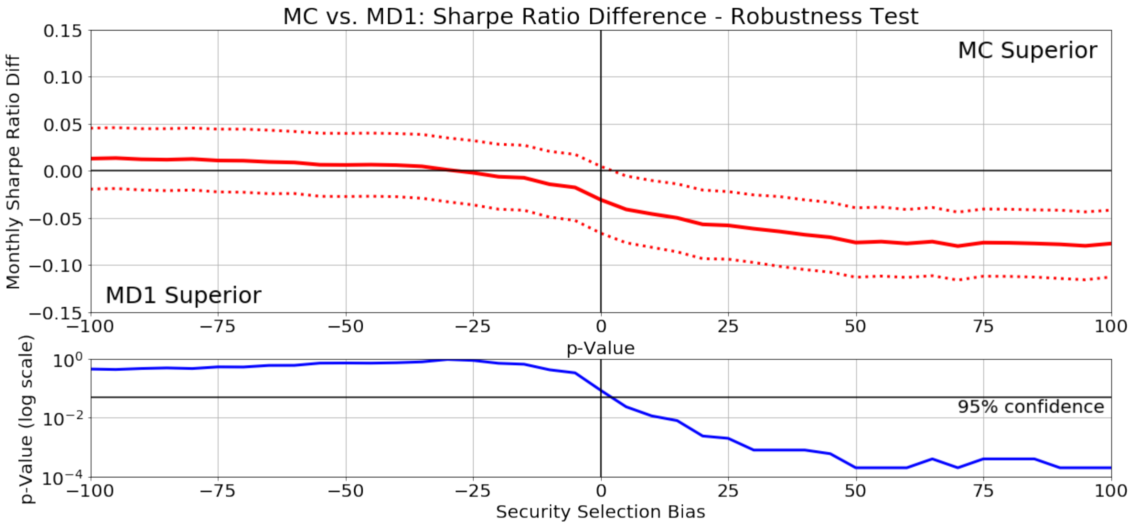

Figure 17: Sharpe ratio contour: MC vs. MD1

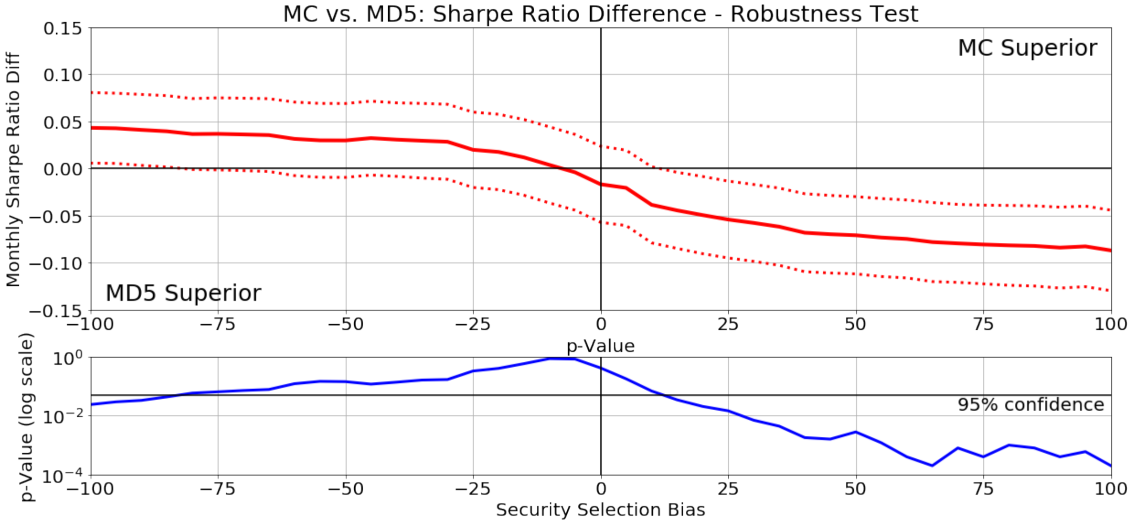

Figure 18: Sharpe ratio contour: MC vs. MD5

3.9 Equal Weight vs. Panel of Allocation Methods

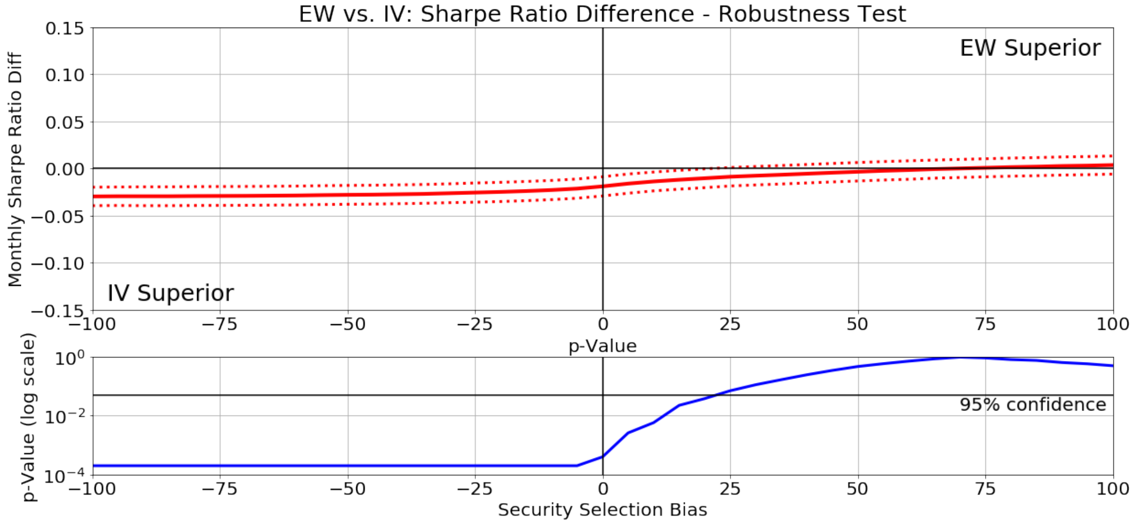

Figure 19: Sharpe ratio contour: EW vs. IV

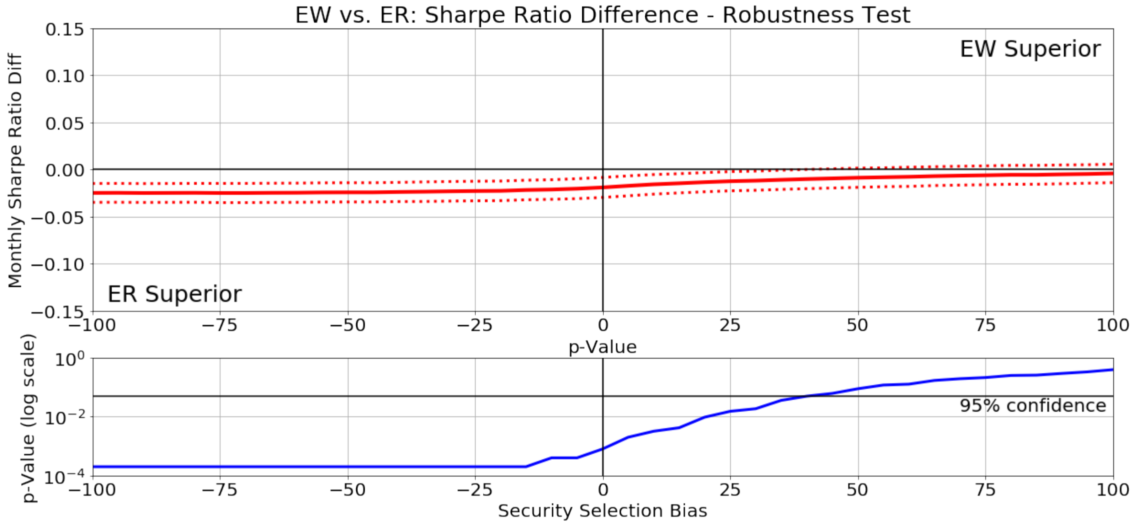

Figure 20: Sharpe ratio contour: EW vs. ER

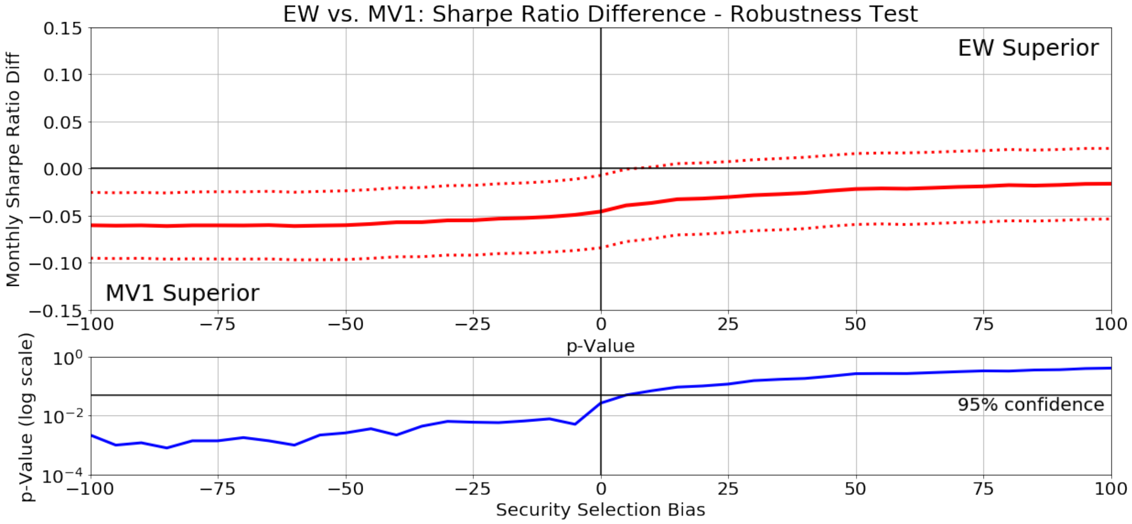

Figure 21: Sharpe ratio contour: EW vs. MV1

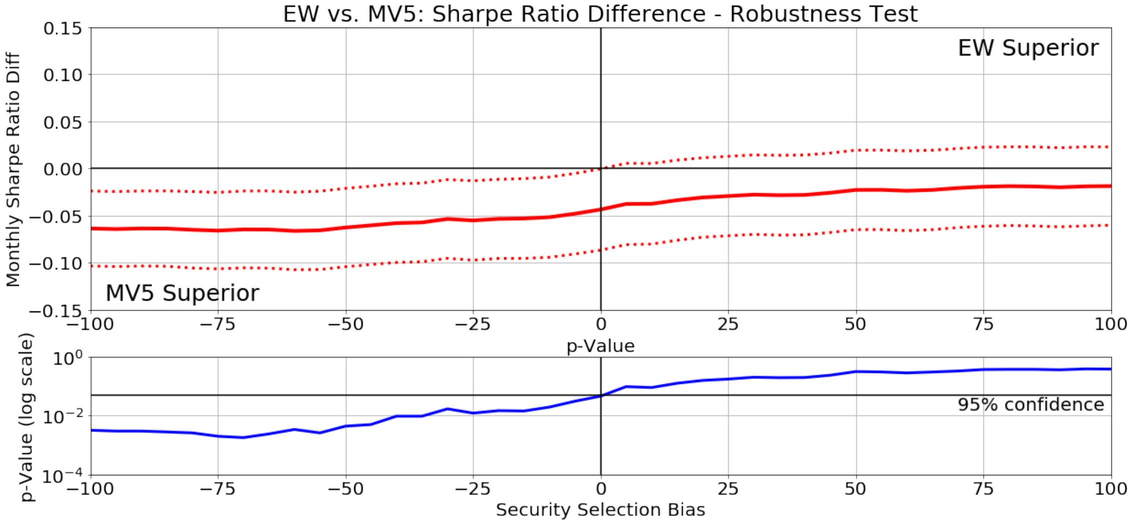

Figure 22: Sharpe ratio contour: EW vs. MV5

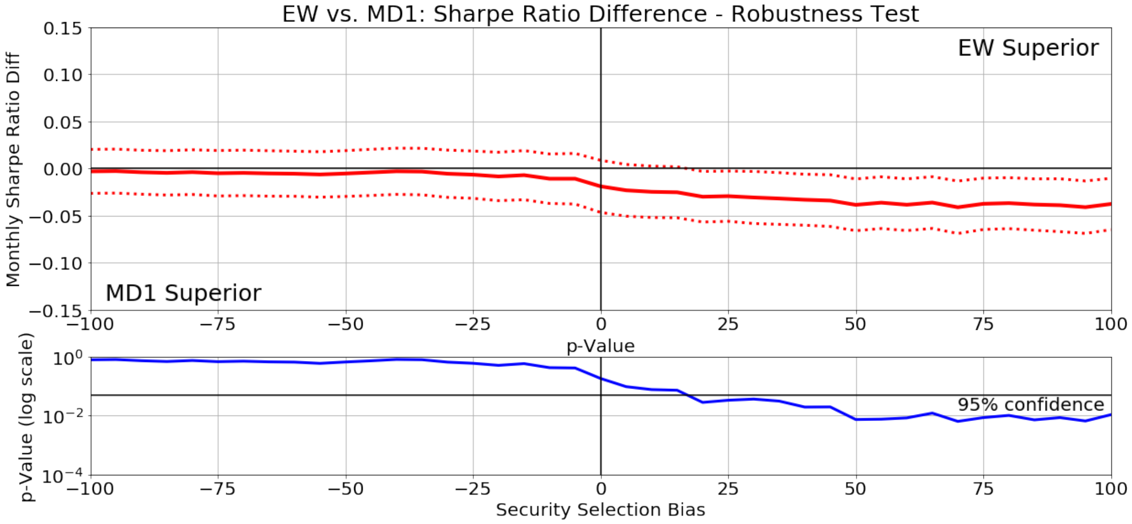

Figure 23: Sharpe ratio contour: EW vs. MD1

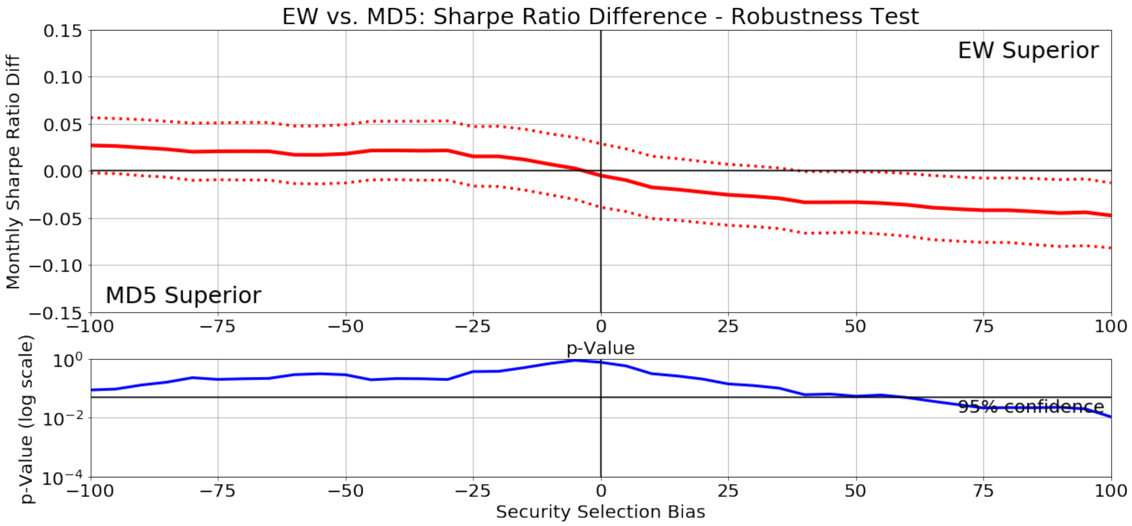

Figure 24: Sharpe ratio contour: EW vs. MD5

3.10 Inverse Volatility vs. Panel of Allocation Methods

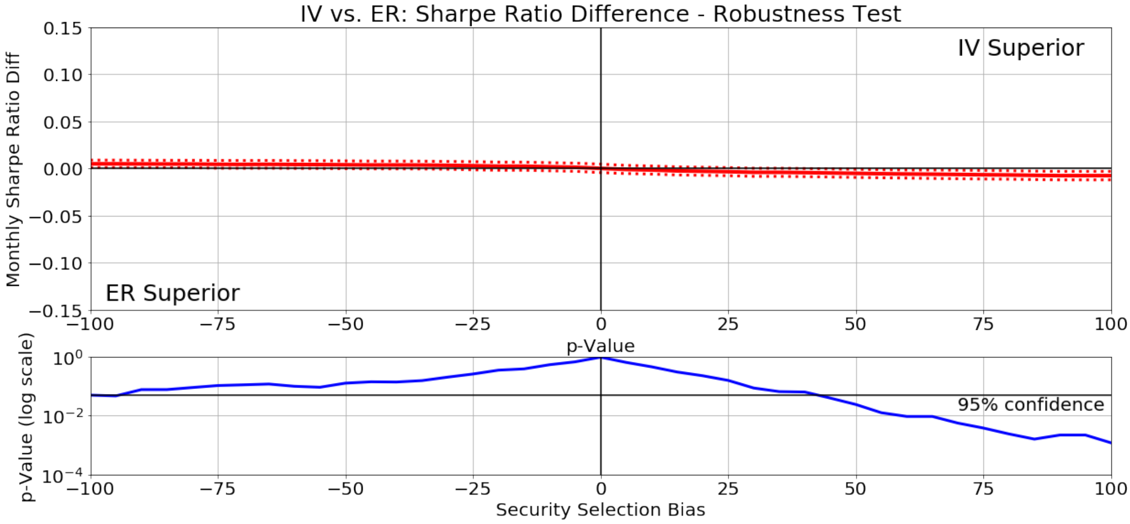

Figure 25: Sharpe ratio contour: IV vs. ER

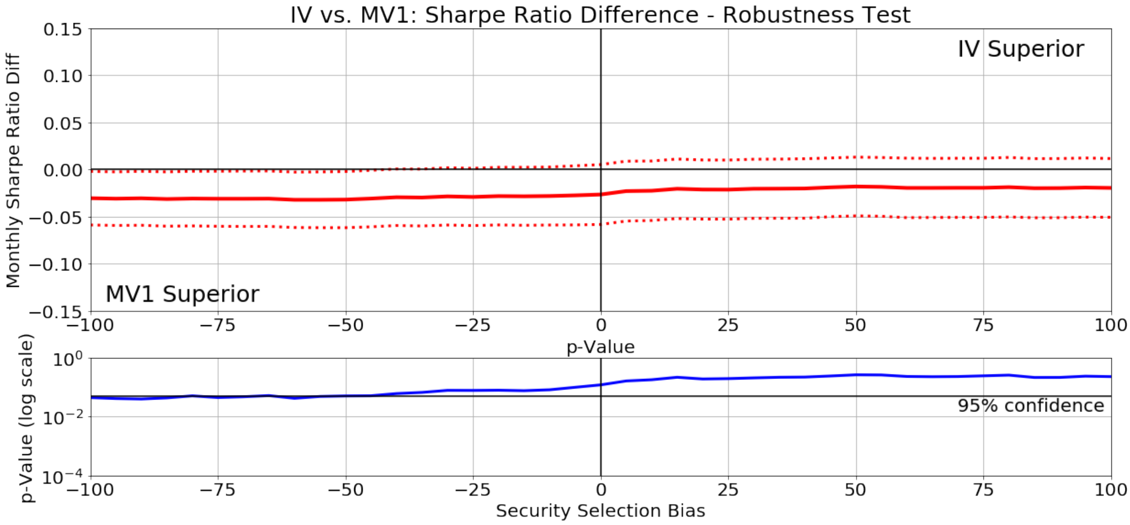

Figure 26: Sharpe ratio contour: IV vs. MV1

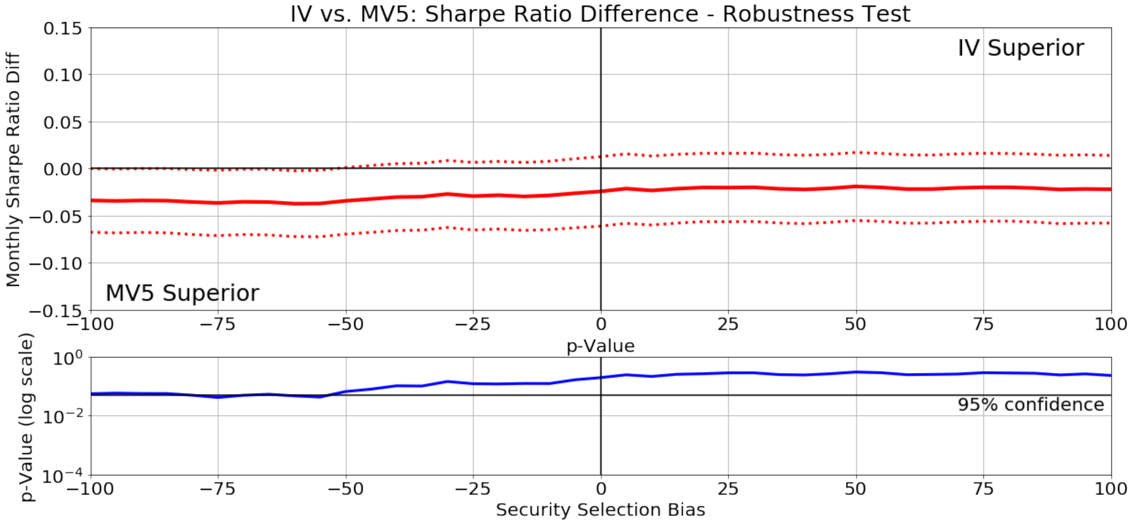

Figure 27: Sharpe ratio contour: IV vs. MV5

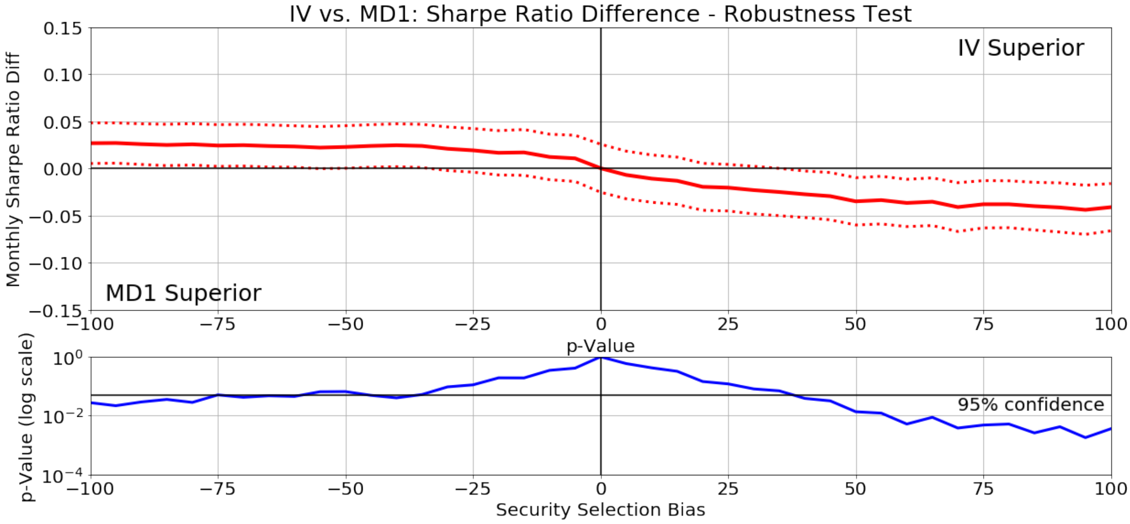

Figure 28: Sharpe ratio contour: IV vs. MD1

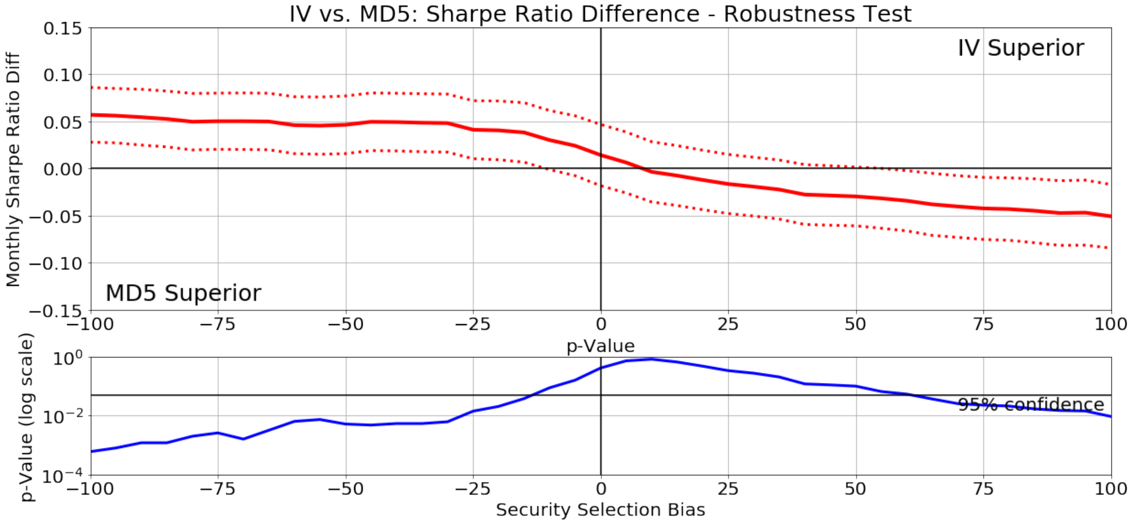

Figure 29: Sharpe ratio contour: IV vs. MD5

3.11 Equal Risk vs. Panel of Allocation Methods

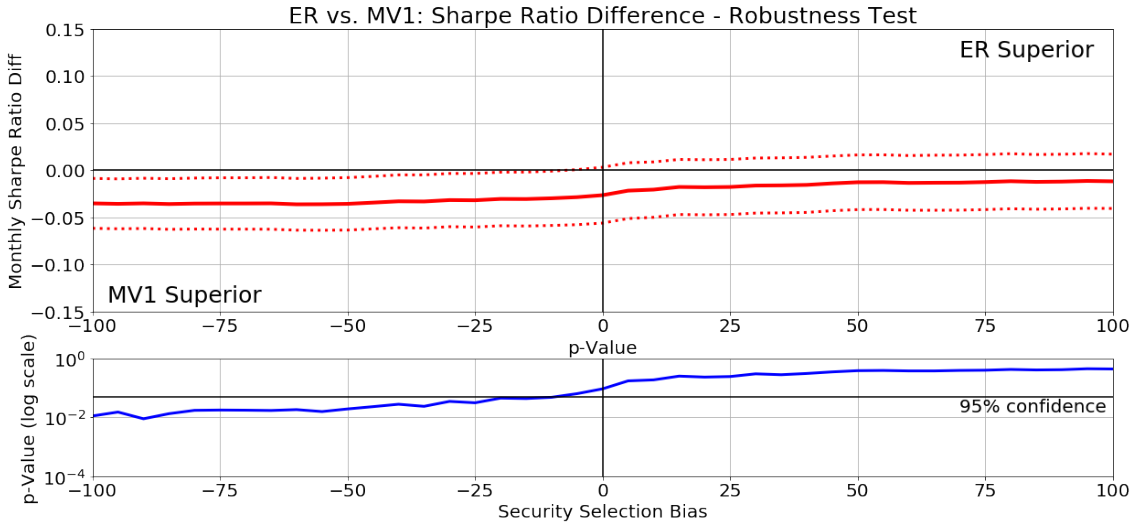

Figure 30: Sharpe ratio contour: ER vs. MV1

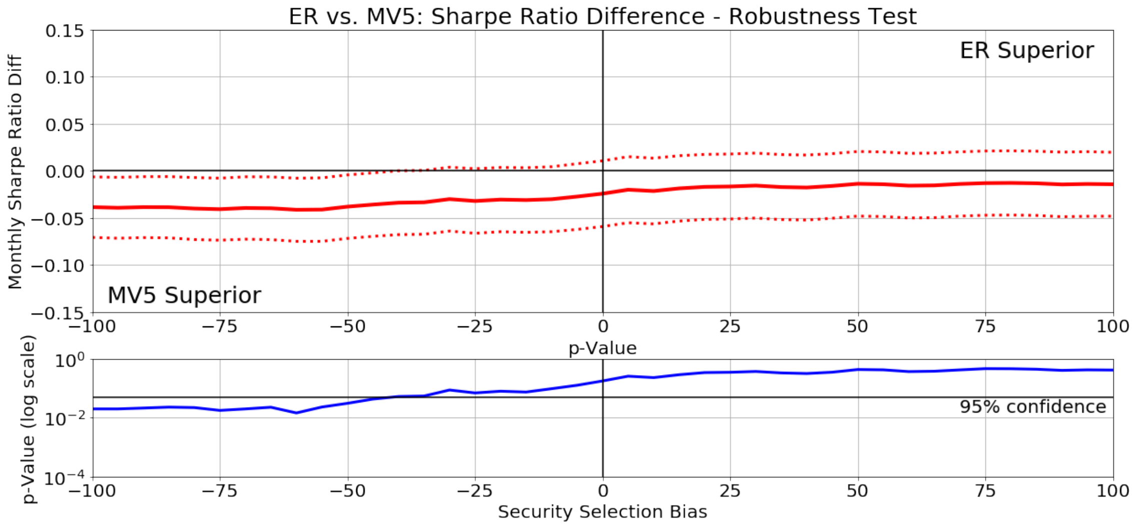

Figure 31: Sharpe ratio contour: ER vs. MV5

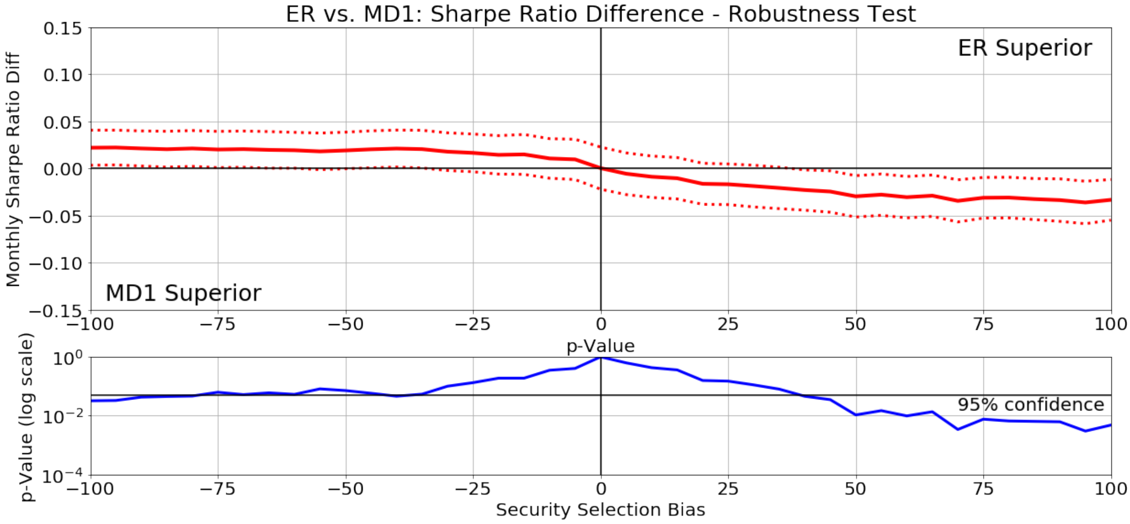

Figure 32: Sharpe ratio contour: ER vs. MD1

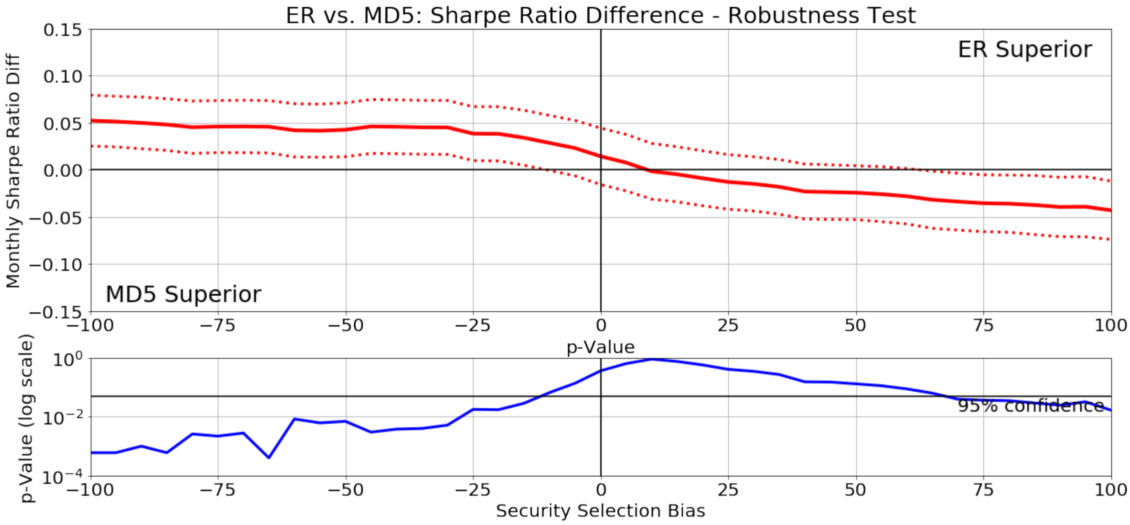

Figure 33: Sharpe ratio contour: ER vs. MD5

3.12 Minimum Variance 1% vs. Panel of Allocation Methods

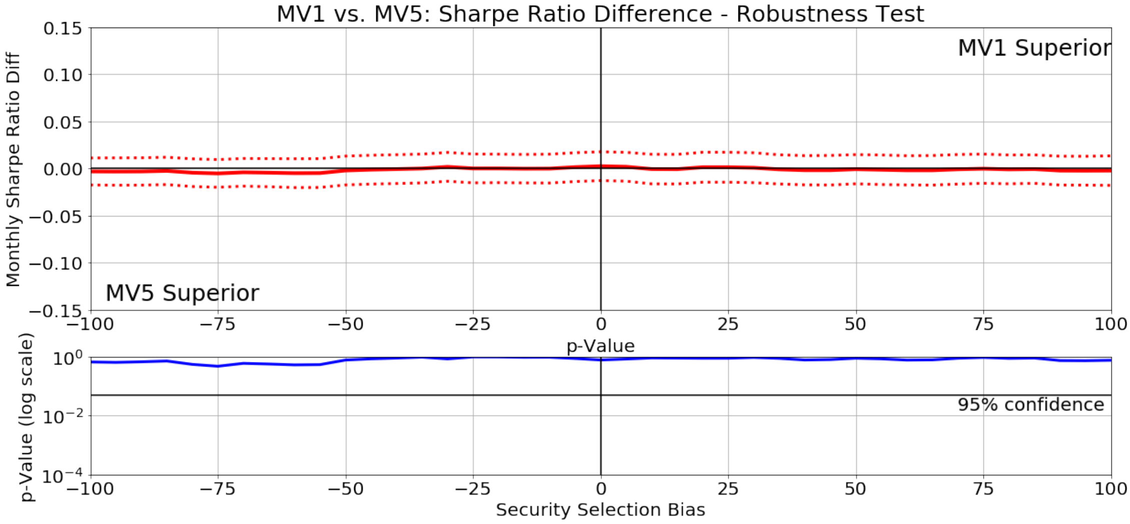

Figure 34: Sharpe ratio contour: MV1 vs. MV5

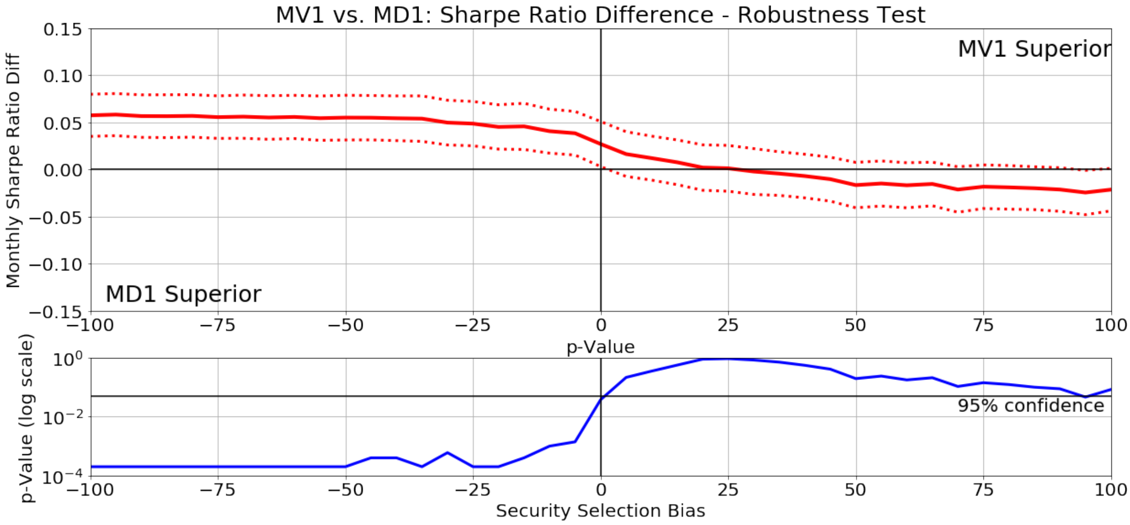

Figure 35: Sharpe ratio contour: MV1 vs. MD1

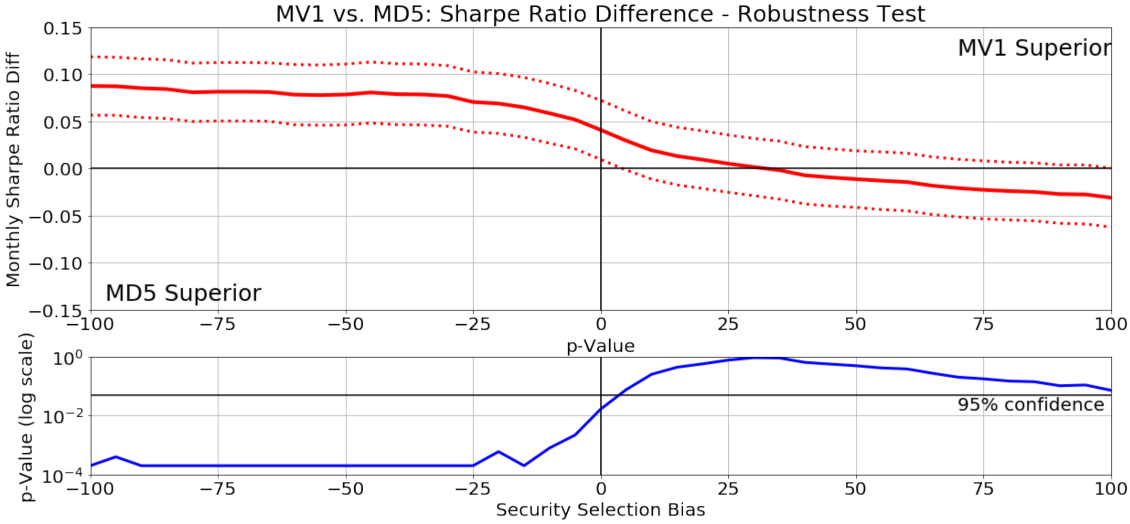

Figure 36: Sharpe ratio contour: MV1 vs. MD5

3.13 Minimum Variance 5% vs. Panel of Allocation Methods

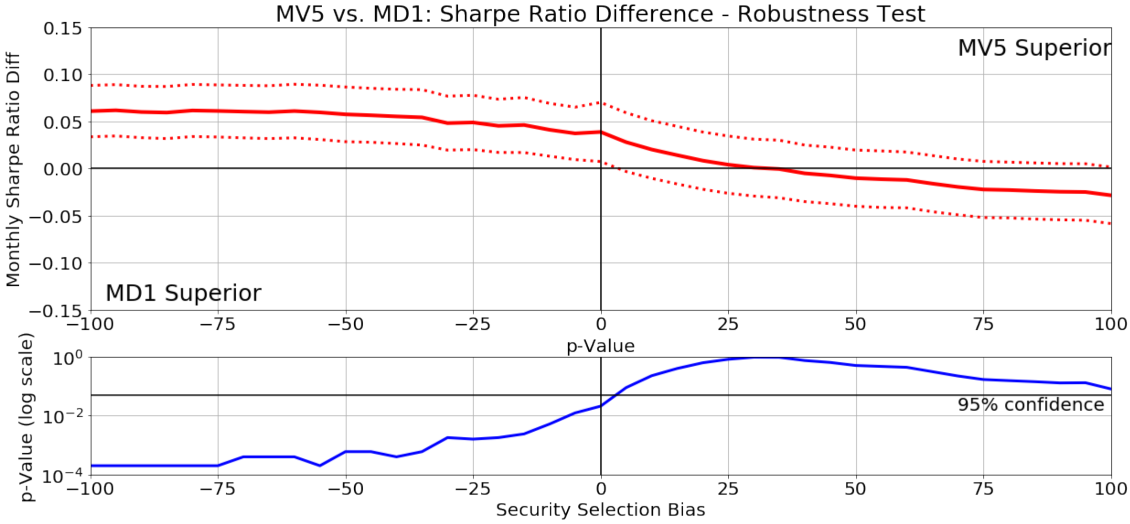

Figure 37: Sharpe ratio contour: MV5 vs. MD1

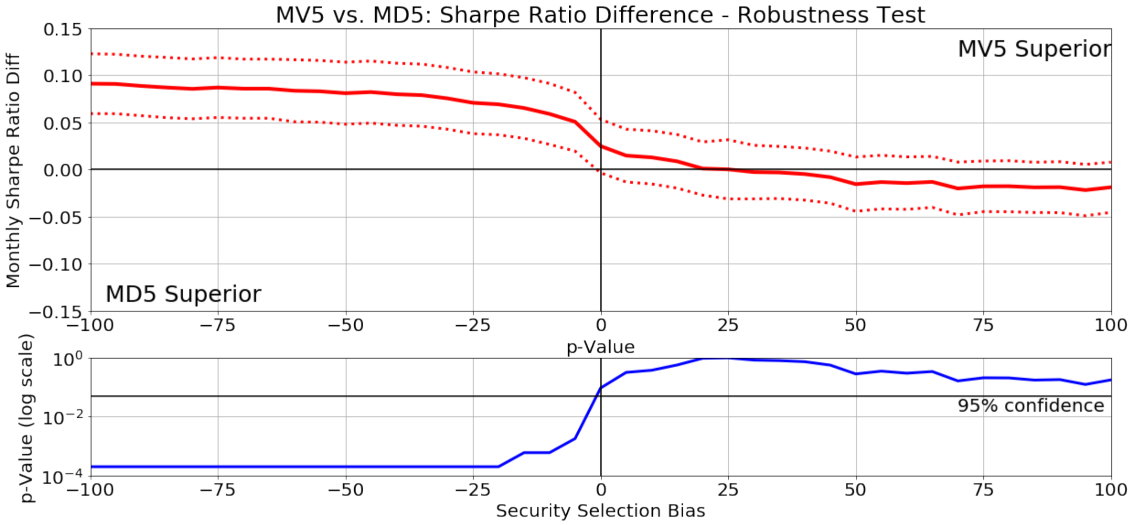

Figure 38: Sharpe ratio contour: MV5 vs. MD5

3.14 Maximum Diversification 1% vs. Panel of Allocation Methods

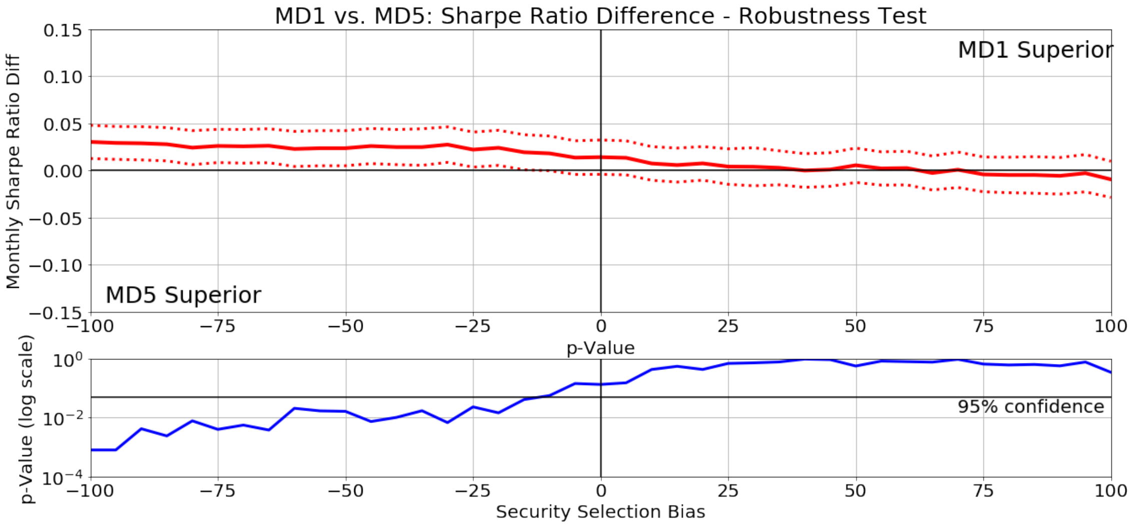

Figure 39: Sharpe ratio contour: MD1 vs. MD5

Chapter 4

Low Correlations Are The Dominant Feature Of Optimal Sharpe Ratio Portfolios

4 Low Correlations Are The Dominant Feature Of Optimal Sharpe Ratio Portfolios

4.1 Chapter Introduction

The three inputs to the Sharpe ratio are portfolio returns, the risk free rate, and the standard deviation of the portfolio returns (volatility). Going more deeply into the formula, portfolio returns are formed from individual stock returns, correlations, and the capital weights applied to the stocks.

This chapter highlights the significance of the correlation structure when configuring portfolios with optimal Sharpe ratios. This finding stands in contrast with what appears to be significantly more practitioner effort focused on predicting security returns when attempting to build optimal portfolios.

In the past, analysts pursued security selection through the fundamental analysis of securities by analyzing spreadsheets, visiting companies, and other approaches such as talking to management and suppliers. Today, the analysis for security selection has advanced through the use of non-traditional methods which include big data, artificial intelligence, machine learning, and the mining of alternative data sets such as satellite imagery, mobility tracking data, and social media.

The finding of this chapter indicates that there is cause for at least an equal amount of effort to be dedicated toward the discovery of methods to optimize the correlations between securities rather than on the relative returns between securities.

This chapter demonstrates that portfolios with optimal Sharpe ratios are dominated by harvesting low correlations, not by selecting stocks with high returns.

4.2 Optimal Sharpe Ratio Portfolio - A Model

4.2.1 Model Purpose

To study the effect of each input to the Sharpe ratio, it is necessary to isolate the individual impacts of returns, volatilities, and correlations in an optimal setting. To do this, a theoretical construct is introduced which is not practical for implementation, but is simple and used as a tool only for analysis. The portfolio construction method is called the Optimal Sharpe ratio portfolio (OSR).

4.2.2 Near-Perfect Foresight

In the previous chapter, the concept of near-perfect foresight was introduced. For the purpose of this chapter, the near-perfect foresight method of security selection is extended to assist with the OSR portfolio formation. A description of the security selection method is also warranted here.

If the objective of portfolio formation were simply to achieve the highest returns possible, irrespective of risk, then perfect foresight would select the one stock with the highest returns. However, Markowitz clearly rules this out as a practical path and sets in place the superiority of diversification.

Diversification is only possible with more than one security. This analysis seeks to understand how the Sharpe ratios from a panel of allocation methods change in the presence of security selection (positive and negative). So to form portfolios from sets of securities with better and worse performance, we need a method to pick stocks each year to feed into the allocation methods and observe the resulting Sharpe ratios.

As the existing literature points out, there is a vast array of methods espoused for selecting stocks. These can be from any combination of discretionary or systematic, and from fundamental and quantitative. Each of these methods can select from a wide variety of inputs (financial statement ratios, relative momentum, etc.) and parameters (1 year look-back, rolling 3 years, etc.). Any method that we might use in our analysis can be adapted with a slight change in parameter set or inclusion of an additional data point. And any method that we use may work well in a particular year and poorly in a later year.

For this reason, I introduce the concept of near-perfect foresight. It is not achievable in practice, but it is simple to implement, reliably creates the stock return bias that is needed for answering my particular questions, and leaves intact the requirement of diversification that Markowitz established. Near-perfect foresight is a boundary condition method in that it is the best outcome that can be achieved while preserving diversification. No other method of security selection could outperform near-perfect foresight and retain the same potential for diversification.

So if perfect foresight implies picking one single best winner, then my versions of near-perfect foresight would be identifying in advance a range of stocks with the biggest gains or losses in the following year and removing them from the set of available securities. This method leaves intact the potential for diversification by continuing to work with a large set of available securities while reliably improving or degrading the return profile of the potential portfolios.

The approach is implemented by using a one year look ahead when forming the portfolios. As before, it begins with the 1,000 largest capitalization stocks available in the data set at the end of the year. I then look ahead a full year and rank the performance of each stock that is available for the portfolios. Portfolios are formed with advance knowledge by systematically removing the best or worst performing stocks based on performance in the following year.

For this chapter, the original near-perfect foresight approach used in the last chapter is now extended to include a one year look-ahead into the additional characteristics of volatility and correlation.

4.2.3 OSR Model Definition

The OSR may seem similar in name to an existing systematic portfolio construction method called the Maximum Sharpe ratio (MSR). The OSR method differs significantly from the MSR. A portfolio constructed using MSR uses historical data to estimate the returns and covariance matrix. The weights are the solution which provides the maximum Sharpe ratio portfolio given the estimated parameters.

Rather than using historical information to estimate parameters and then to optimize portfolio weights, the OSR method introduced here is forward looking and based on the same concepts from the near-perfect foresight used in the previous chapter. The OSR method creates the most efficient portfolio that could have been constructed with one-year perfect foresight into returns, volatilities, and the correlation matrix.

If the parameter estimates used for the MSR portfolio turned out to be perfect, then the MSR portfolio would equal the OSR portfolio.

To find the weights for the OSR portfolio, the maximization requirement is given by:

$$\underset{w}{arg\,max}\,\frac{\sum_{i=1}^{N}[w_i r_i] - R_f}{\sqrt{w'\sigma w}}$$

The result of the maximization routine produces the vector of weights for each stock in the portfolio with w_i being the weight of each asset, r_i being the return of each asset, R_f being the risk free rate, and $\sqrt{w'\sigma_i w}$ being the standard deviation of the portfolio.

The notations of w and w' are the vectors of the weights from the resulting maximization.

Some of the allocation methods used in my previous chapters have required an optimization using the covariance matrix which is formed from historical estimates. Because optimizers amplify the impact of covariance outliers, the matrix is normally adjusted according to the shrinkage routine of (Ledoit and Wolf 2004).

Quite important to note however, is that in the case of the OSR portfolio, I do not use the shrinkage method for the covariance matrix because I am not estimating this parameter but instead am looking ahead one year with perfect foresight.

4.2.4 OSR Portfolio Formation

Before beginning to focus on the performance characteristics of the OSR, it is important to examine the consequences of the allocation method on the size of the security set selected. While the overall method is not achievable without perfect foresight, the portfolios resulting from the maximization routine should not be overly concentrated in a small set of securities.

In the near-perfect foresight portfolios used in Chapter 3, I restricted the method to removing a small subset of the best or worst performing securities to create a reasonable return tilt in the universe of available stock returns. The tilt was entirely based on foresight into returns, not volatilities or correlations. The method used in Chapter 3 of removing a few stocks preserved a large portion of the available securities so that diversification remained possible. Had full perfect foresight been used, an unreasonably small set of securities would have been used in the formation of the portfolios. Perfect foresight in our context would have started with batches of the 5 best or worst securities. This is too small of a subset to be reasonably considered.

By contrast, the OSR method uses full perfect foresight but seeks a more complex objective. Because the optimization functions on a ratio of portfolio returns and standard deviation, and because there are low and stable correlations in the cross section of returns, a reasonable number of securities are selected each year as shown in Figure 40.

The OSR method imposes two important restrictions on the weights of individual securities. First, so as to maintain a comparison with the other allocation methods studied in our panel, a maximum security size of 5% is allowed. Also, so as to reduce the number of insignificant positions, a minimum position size of 0.01% (1 basis point) is imposed.

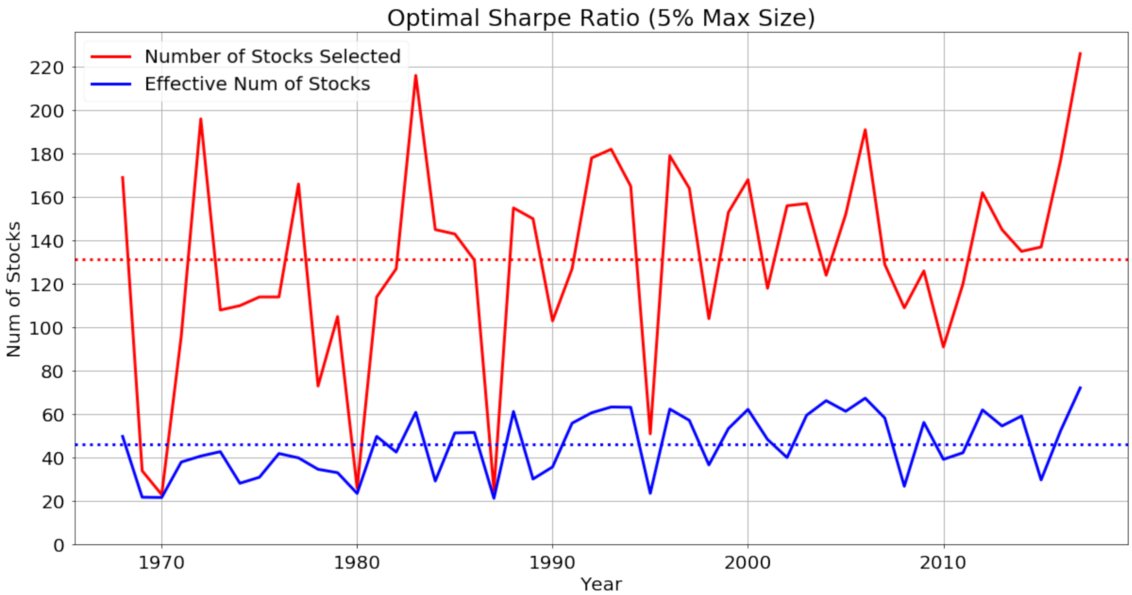

Figure 40: OSR Portfolio: Number of Stocks Selected
The figure above displays the number of stocks selected from the top 1,000 market capitalization universe each year to form the optimum Sharpe ratio (OSR) portfolio with perfect foresight. The mean observation across the time series is 131. By accounting for the concentration of security weights in the OSR, the figure also displays the effective number of stocks in the OSR, with a mean observation of 46.

The number of stocks selected is a simple measure of the starting point each year. However, the Effective Number of Stocks (described in Appendix 1) measures the concentration impact of the weighting that the OSR portfolio ultimately selects. It is this measure of concentration which can be compared to prior academic studies to check the reasonableness of the concentration of the OSR portfolios.

4.2.5 Concentration - Academic & Practical Reference Points

There is a section of the literature which relates portfolio concentration (measured in various ways) to the observed performance. But summary data on what comprises a concentrated fund is limited. This is due in part to the discretionary nature of how to classify a concentrated portfolio. However, a few academic studies have shown data on concentrated manager cohorts which can be referenced to check the reasonableness of the portfolio sizes.

Even though a stock set of 131 is significantly smaller than the 1,000 stocks that I start with each period, it is in line with the findings of (Fulkerson and Riley 2019) who study portfolio concentration from 1999 through 2014. In their study, the median number of stocks is around 180 while the first quartile by selection size is around 140. As will be shown later, the starting point for my analysis will be an equally weighted portfolio of the most optimal stocks. Equal weight portfolios have an effective number of stocks equal to the total number of stocks as there is no concentration effect. This connects the work of (Fulkerson and Riley 2019) to my starting point as a reasonable reference.

In the years 1970, 1980, and 1987, the OSR process produces the minimum stock selection sets of 23, 26, and 25. While these are the lowest levels observed in the OSR data set, it is within the observed range of securities used by concentrated active equity managers. This range of holdings is supported in research by (Cremers and Pareek 2016), a study that spanned both US equity mutual funds as well as institutional portfolio holdings from 13-F filing. Of 39,555 portfolios examined, the mean holding size was 117 with a standard deviation of 195. The smallest portfolio held just 10 stocks.

In practice, offerings in the marketplace are observed that advertise high levels of stock concentration. A search of Morningstar for concentrated managers shows top results as:

1. Alliance Bernstein Concentrated Growth Fund. Ticker (WPASX). "Approximately 20 primarily US large-cap stocks". https://www.alliancebernstein.com/funds/us/equities/us/growth/concentrated-growth.htm

2. Lazard US Equity Concentrated. Ticker (LEVIX). "The strategy typically invests in 15–35 companies with market capitalizations generally greater than $350 million." https://www.lazardassetmanagement.com/us/en_us/investments/strategy/us-equity-concentrated-strategy/s81

After referencing both the academic studies and a review of market offerings, the concentration of the OSR portfolios in any year seems reasonable.

4.2.6 OSR Model Design Summary

The OSR portfolio is clearly not achievable without perfect foresight. However, just as with the near-perfect foresight method crafted for security selection, the OSR perfect foresight method is created as a boundary condition for the purpose of studying and quantifying the potential contributions available from security selection and allocation method.

4.3 Performance Attribution of the OSR Portfolio

This section will examine the performance of the OSR method and illustrate how low correlations are the key driver of portfolios with optimal Sharpe ratios.

4.3.1 Comparison Portfolios

It is accepted that the OSR is not an achievable construction method due to lack of perfect foresight, but is used as a tool to examine the contribution to performance improvements that are possible from varying the available securities and adjusting the capital weighting methods.

To explore these parameters in greater detail, the full perfect foresight OSR method can be compared with three other portfolio allocation methods.

First, the opposite of perfect foresight is no foresight. The equal weight (EW) method using the full 1,000 largest capitalization stocks is the embodiment of no foresight.

Second, perfect security return selection is examined. This method focuses purely on selecting the highest returning securities without regard for volatility or correlations. For the comparison in this section of analysis, we will assume perfect foresight with regard to security performance and select the 20 best performing stocks each year. Since the OSR method is formed with a max position of 5% in any single stock, then a set of the 20 best performing stocks has the greatest level of concentration that would be allowable under the OSR method. This makes the two methods comparable in potential concentration.

For the final comparison method, the stocks selected for the OSR portfolio are equally weighted, not optimally weighted. The full OSR method knows all of the returns, volatilities, and correlations for the full set of 1,000 stocks a full year in advance. The OSR method both selects the optimal security set, and applies the perfect weighting scheme to maximize the Sharpe ratio. For the final comparison, the optimal weights are omitted and only an equal weight portfolio of the OSR securities is shown. By equal weighting the OSR selected stocks, we can isolate the impact of security selection in the OSR optimization.

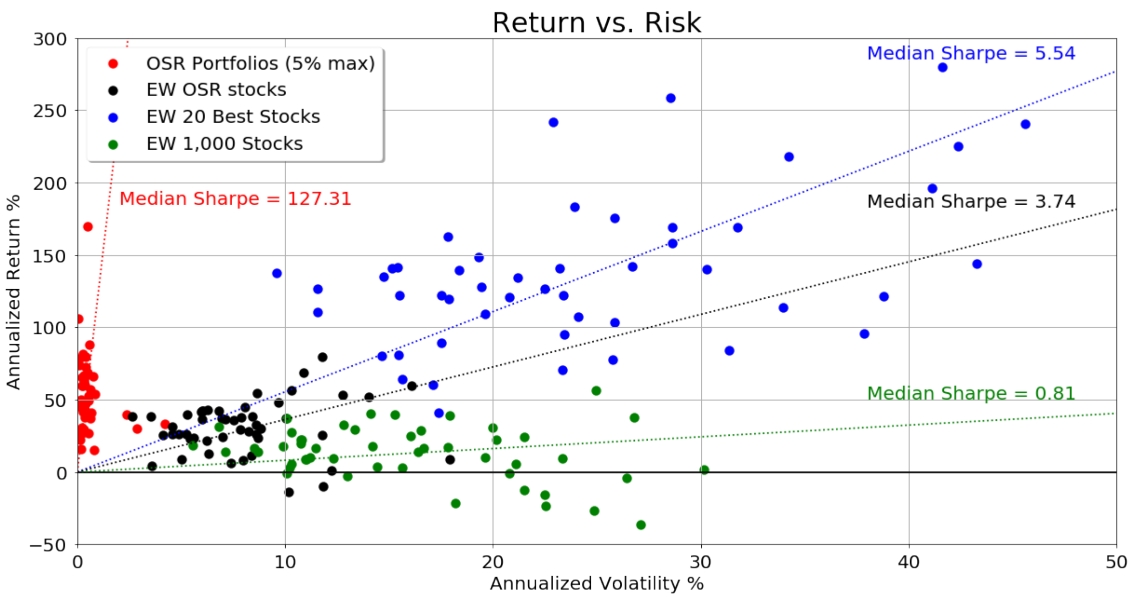

Figure 41: OSR Performance Comparison
The figure above displays the annual return vs. risk for each allocation method through the 50 years of the study, 1968 to 2017. Each dot represents a single year within the study period for each method of portfolio construction.

4.3.2 Initial Performance Observations

The median Sharpe ratio of the EW 1,000 Stocks (green) is in line with the previous summaries of performance that have been shown in the opening chapters and in the data summary in the appendix. This is the baseline for the comparisons as there is no security selection, no optimal weighting, and no foresight.

The median Sharpe ratio of the portfolios which select the 20 best performing stocks each year (blue), is significantly above the EW portfolios. This is not much of a surprise, however, my expectation was for these observations to be substantially higher than 5.54. Annualized returns are significantly higher, but annualized volatility also increases which reduces the observed Sharpe ratios. The ability to perfectly pick the top 20 stocks each year is no doubt an advantage (even if unattainable) over the EW method, and accounts for a 4.73 points increase in median Sharpe ratio.

The OSR method (red) with full perfect foresight exhibits median Sharpe ratios over 127. Again, these are clearly unattainable in practice, but it is curious how these enormous gains in efficiency are derived. The red dots show improvements in annual returns over the EW (green), but the main source of improvement is the enormous reduction in portfolio volatility. This can be seen by the clustering of the red dots around the zero level of annualized volatility.

The clustering is a function of volatility and correlation of the individual stocks, but it is not yet clear from the portfolio performance which parameter is the key driver of the optimal Sharpe ratios.

Finally, EW of the OSR stocks (black dots) are the same stocks that are used to construct the OSR (red) portfolios. By equal weighting the OSR stocks, the performance aspect can be isolated from the weighting aspect. From this we see that the black dots are slightly better in terms of annual performance than the green dots (no selection, EW), but are significantly less volatile even with an equal weighting.

The black dot portfolios are clearly not picking the highest returning stocks despite having perfect foresight into performance. These portfolios are being formed from stocks with slightly better performance than average, but with some combination of lower volatility stocks with lower average correlations. This means that the enormous improvement in the Sharpe ratios of the red portfolios is being harvested primarily from the volatility and correlation combinations within the security set. Harvesting these characteristics is optimized through perfect weightings.

4.3.3 Time Series Isolation

Now we move from portfolio results to examining the inputs of the underlying stocks. To see the comparison of the return, volatility, and correlation elements in more detail, the full study scatter plot approach is now shown in time series format. For each year, the mean observation of the stocks in each portfolio is shown for each parameter. This allows an annual comparison as another way to see how each of the annual dots match up with each other.

The three methods shown in the following charts all use equal weights of the securities which levels the field for a fair comparisons. The OSR (red) portfolios are not included as they are just a special optimal weighting of the black dot portfolios.

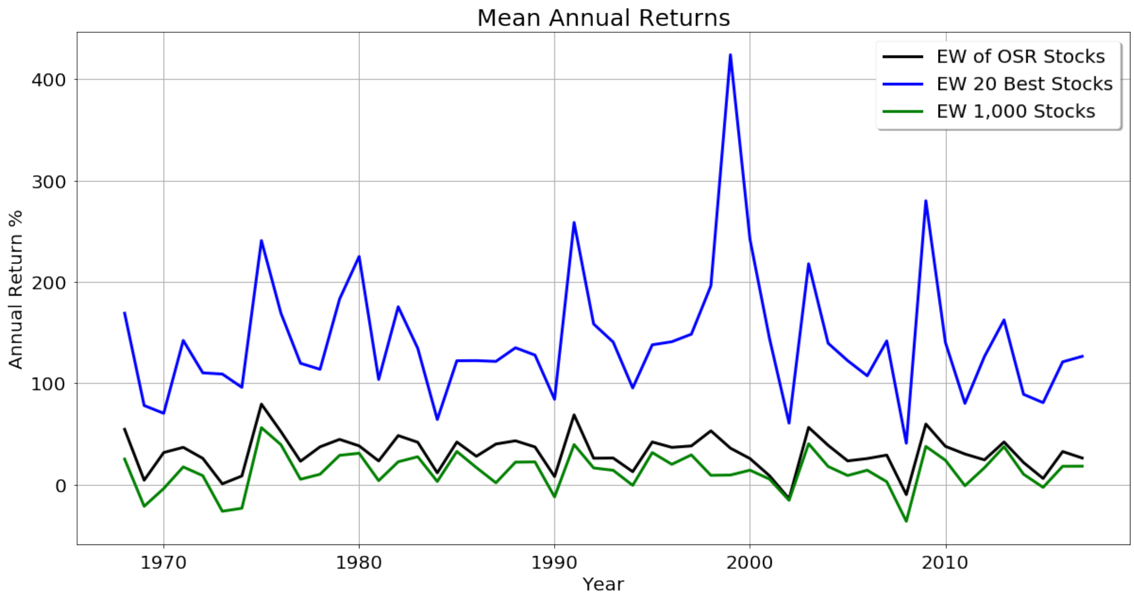

Figure 42: Average Security Performance Comparison
The figure above displays the annual mean return of the securities selected for inclusion in the allocation methods.

4.3.3.1 Mean Annual Returns In Figure 42, the scatter plot of Figure 41 comes to life year by year with observations of the mean annual returns of the component stocks. Each of the methods selects securities that have better mean performance than the full set of 1,000 stocks. The 20 best performing securities (blue) are significantly above the full 1,000, as expected. These portfolios are formed from the very top of the range of security performance available each year.

The EW of OSR (black) portfolios at first appear to select only marginally better securities than the full 1,000 securities available each year (green). However, on closer inspection, the vertical scale of the graphic makes it difficult to see that the mean performance of the EW OSR securities is still meaningfully above the EW 1,000 stocks.

It is now clear that the large improvement in the OSR portfolios are not being derived primarily from perfect foresight into security performance. This evidence stands in contrast to the effort being expended in the industry to select securities with superior performance.

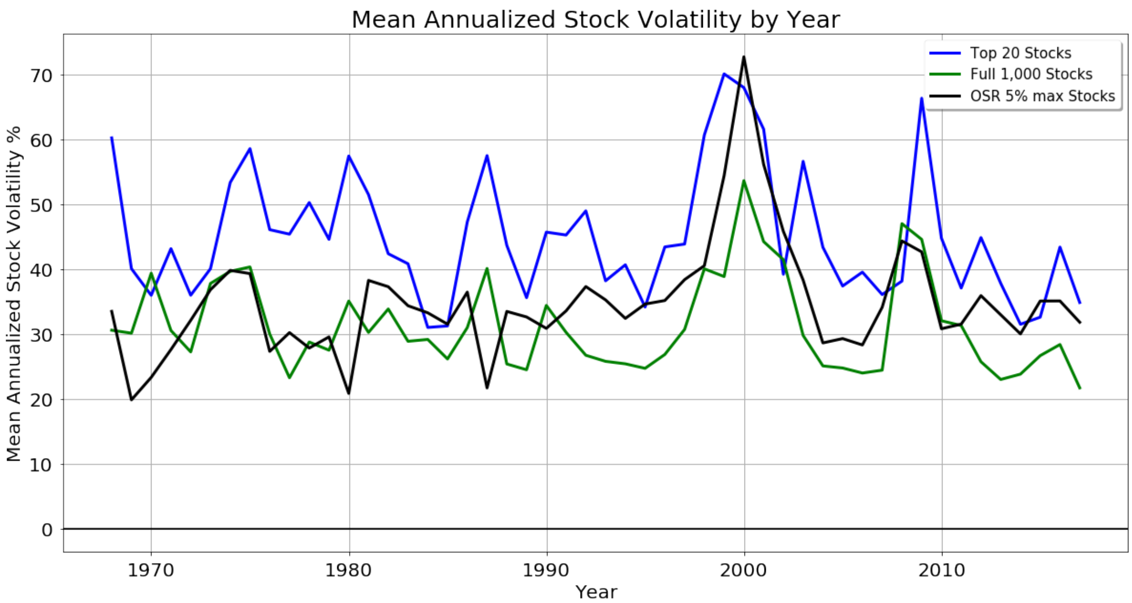

Figure 43: Average Security Volatility Comparison
The figure above displays the mean annualized volatility each year of the securities selected for inclusion in the allocation methods.

4.3.3.2 Mean Annual Volatility

The mean volatility of the stocks in each of the allocation methods is more similar than the returns that we saw in Figure 42. The higher returns of the top 20 stocks are accompanied by higher volatility throughout the time series. The mean volatility of the OSR stocks is very similar to the mean volatility of the full 1,000 stocks available each year in our study.

This observation further explains that volatility differences are not the driver of the large Sharpe ratio improvements. All else being equal, lower volatility securities could lower the volatility of the portfolio (denominator of the Sharpe ratio) and increase the Sharpe ratio. But this is not what is happening with OSR portfolios.

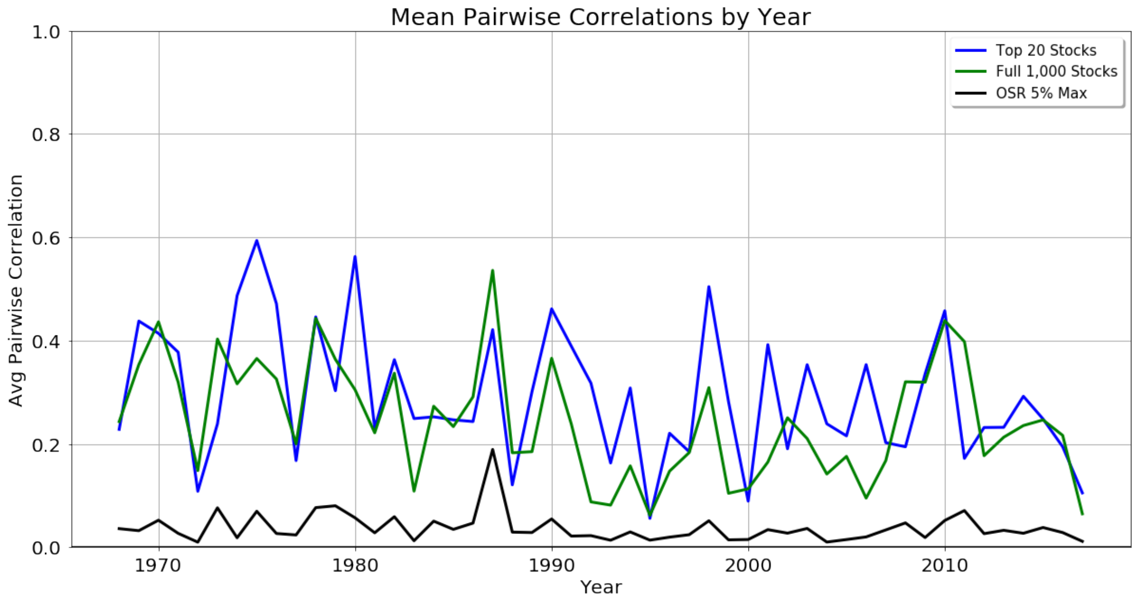

Figure 44: Average Pairwise Correlation Comparison
The figure above displays the mean pairwise correlation of the securities selected for inclusion in the allocation methods.

4.3.3.3 Mean Pairwise Correlations

It is quite clear from Figure 44 that the OSR method is selecting securities with dramatically lower mean pairwise correlation. The mean pairwise correlations of the top 20 performing stocks are indistinguishable from those of the full set of 1,000 stocks.

Looking at the time series of mean annual input characteristics makes clear that the OSR portfolios are constructed from optimal weighting being applied to a set of securities with the richest set of diversification potential.

4.3.4 Security Selection Crowds Out Diversification Benefits

A portfolio formation process that starts with security selection based on predicted returns directly limits the potential diversification benefits of the portfolio. As portfolios become more concentrated, this potential benefit is severely limited and even perfect selection based on returns can not overcome the cost of limited diversification.

To show that low correlations are much more important than security selection in the creation of optimal Sharpe ratio portfolios, we now look more closely at the top 20 performing securities each year. In our previous analysis shown in Figure 41, the blue dots of the Top 20 performing securities are shown with equal weights. This was done to maximize the contribution possible from security performance selection, while limiting the concentration to the same level as possible with the OSR portfolios.

However, when the concentration constraint of maximum 5% in any security is relaxed, the limiting effect of primarily focusing on stock level performance can be quantified.

Figure 45 starts with the same blue dots from Figure 41. These are equally weighted portfolios of the top 20 performing securities selected each year with perfect return foresight. These portfolios had meaningfully higher Sharpe ratios than EW of the full 1,000 securities, as expected. By comparison, OSR portfolio Sharpe ratios were more than 20 times higher than these Sharpe ratios. One main difference is that OSR portfolios also included optimal weights.

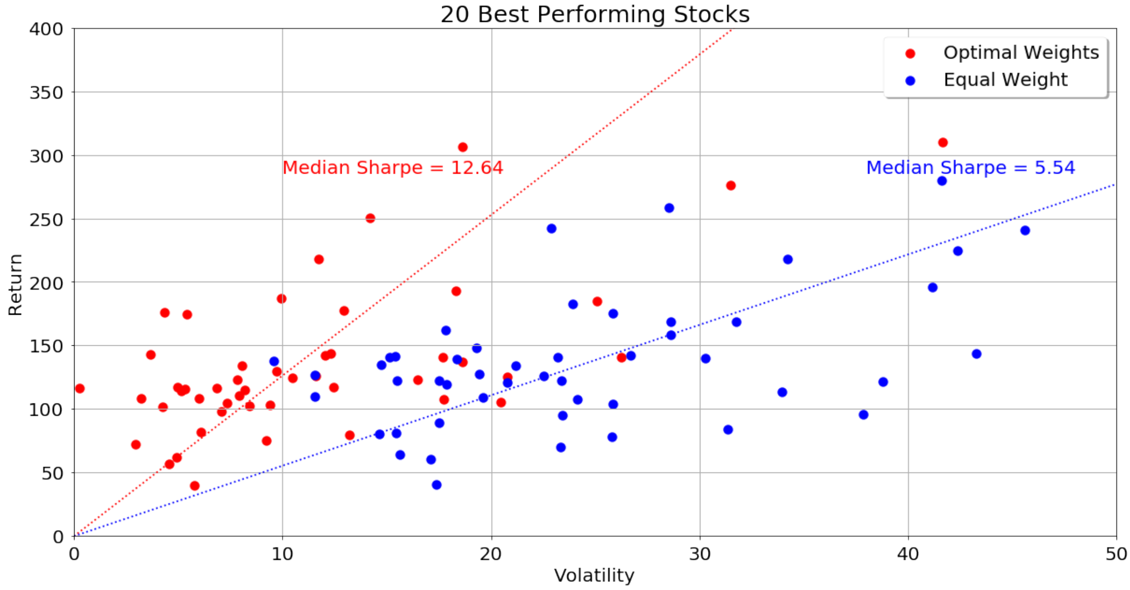

Figure 45: Optimizing the 20 Best Performing Stocks
The figure above displays the Sharpe ratio improvements that are available by optimizing the weights of the 20 best performing stocks each year.

In Figure 45, we start from the same top returning securities used for the blue dots, but then fully relax the weighting concentration limits and optimize the Sharpe ratios by varying the weights.

With free reign to concentrate the portfolio in a very dramatic way, as shown in Figure 46, the median Sharpe ratios improve to 12.64, which is more than double the starting values. While these are incredibly high and unattainable Sharpe ratios, they are still only about 10% of the Sharpe ratios of the OSR portfolios.

This approach started with perfect security return selection as the first step in the formation of the portfolio. This was followed by setting the weighting for each of the stocks in the portfolio. Both of these steps were done with perfect foresight. The net effect of engaging in security return selection first is to limit the potential diversification benefits of the portfolio and reduce the potential Sharpe ratio. Even perfection in return based stock selection can have limiting effects on potential portfolio Sharpe ratios.

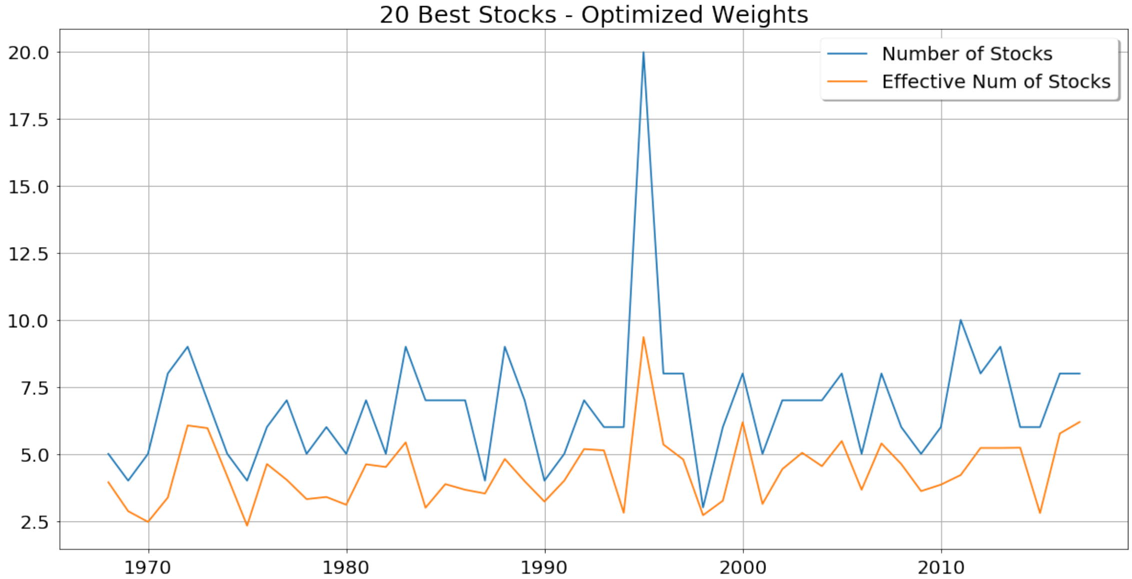

Figure 46: Selection and Weighting of Optimize 20 Best Performing Stocks
The figure above displays the number of stocks selected and the effective number of stocks used when starting from a selection set of the 20 best performing stocks each year, and then allowing the weights to fully float in order to optimize the Sharpe ratio. As in previous optimization methods, the minimum security size allowed is 0.01% (1 basis point). Securities with weights below 1 basis point are set at zero.

Availability of the lowest pair-wise correlations are a necessary condition to creating portfolios with optimal Sharpe ratios. The consequence of engaging in security selection as the primary step to forming a portfolio, is that it creates an inherent limit on the improvement to the Sharpe ratio of any portfolio.

4.4 One Perfect Foresight

4.4.1 Introduction

Selection of securities for portfolios could be made using not only returns, but also volatilities or correlations. As they are all inputs to the Sharpe ratio, each of them have varying effects. But each has a different potential to contribute to an optimal portfolio.

Examining each of these parameters in isolation reveals another important view into the landscape of these characteristics. To expose the landscape of contribution to Sharpe ratio, we begin by forming portfolios with perfect foresight into just one of the parameters.

4.4.2 Method

The landscape of contribution to Sharpe by parameter is evident when we isolate these components and sort them into quartiles.

For the return quartile portfolios, a one year perfect foresight is applied to the full 1,000 stocks each year and stocks are ranked according to their returns. Quartile 4 contains the 250 worst performing stocks each year, while Quartile 1 contains the 250 best performing stocks. The same method is used for forming portfolios based on volatility and correlation where Quartile 4 contain low correlation stocks or stocks with the lowest pairwise correlations. Quartile 1 contain stocks with the highest volatility or highest pairwise correlations.

This method is implemented each of the 50 years in our study and median Sharpe ratios are observed for each perfect foresight parameter and for each quartile. These 12 observations form the rough landscape of perfect foresight by parameter. The landscape will be examined with higher granularity in the next chapter.

4.4.3 Observations

4.4.3.1 Returns In Figure 47, the blue line references portfolios formed with perfect foresight into the returns of the stocks each year. Quartile 4 represents portfolios formed by an equal weight of the bottom performing 250 stocks each year, while Quartile 1

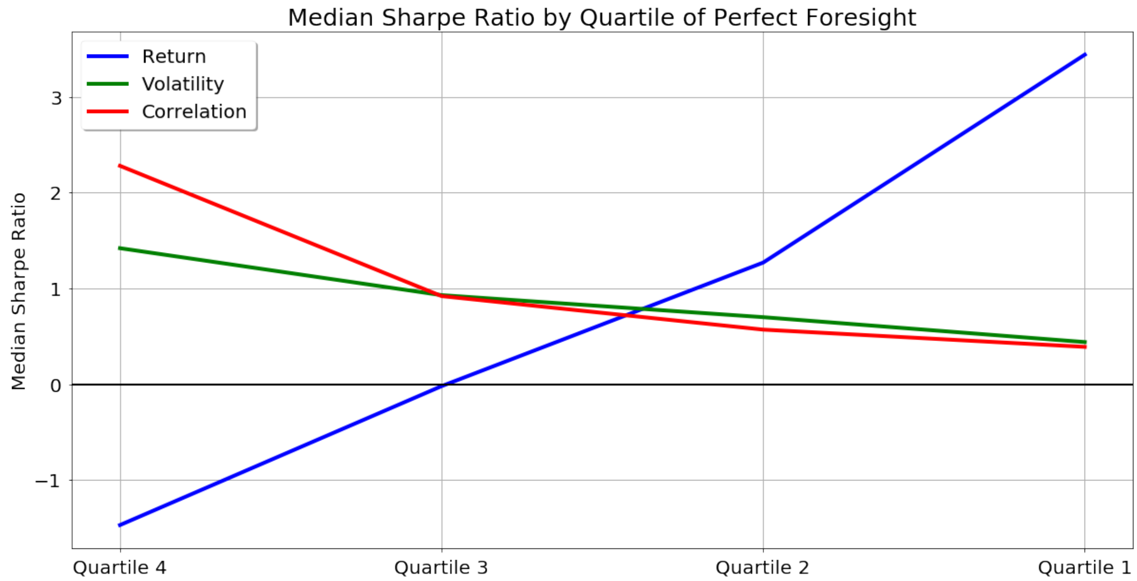

Figure 47: The Landscape of Perfect Foresight by Parameter
The figure above displays the median Sharpe ratio of portfolios formed with perfect foresight into each of the single parameters and sorted by quartile.

represents portfolios form by an equal weight of the top 250 performing stocks each year.

As expected, stocks with higher annual performance naturally form portfolios with higher Sharpe ratios than those with lower performance. Relative to the parameters of volatility and correlation, the slope of the line is steeper indicating a greater reward for selecting well. The consequence of selecting badly is accompanied by a significant penalty.

4.4.3.2 Volatility In the same way, the green line represents portfolios formed with perfect foresight into the volatility of stocks. Quartile 4 represents low volatility stocks while Quartile 1 represents high volatility stocks. Initially it seems perplexing that high risk is not accompanied by higher returns and higher Sharpe ratios. But the slope of this line conforms to the work of (Baker and Haugen 2012) who point out that "The fact that low risk stocks have higher expected returns is a remarkable anomaly in the field of finance." Their work looked at the returns of individual stocks relative to volatility, however, in this analysis, the anomaly manifests itself in the resulting portfolios in the same way.

The low volatility anomaly is exploited by the Minimum Variance portfolio allocation method which showed remarkable robustness in the tests of Chapters 2 and 3.

While there is benefit for correctly selecting low volatility stocks, the penalty for incorrectly selecting high volatility stocks still allows for a positive Sharpe ratio. This is in contrast to the penalty seen for incorrectly selecting securities based on return performance.

4.4.3.3 Correlation In a similar manner, the red line is formed by stocks sorted with perfect foresight into quartiles of pair-wise correlations. Quartile 4 has the lowest pair-wise correlations while Quartile 1 has the highest pair-wise correlations. In the lowest quartile of pair-wise correlation, we also see an acceleration of the benefit to Sharpe ratios.

It is remarkable that the level and slope of this line is quite similar to that of the portfolios sorted into volatility quartiles. There appears to be a low correlation anomaly at the portfolio level, however, these results only apply to a portfolio formed with equal weights and will be examined more closely in the next chapter.

4.4.3.4 Landscape The landscape of Figure 47 contains several interesting observations. First, the return slope is very steep. This may seem to attract the attention of practitioners to predict the best returns. Success analyzing stocks in this manner can lead to relative return performance. However, errors in this approach can just as easily be costly and net out potential gains. The difficulty in predicting the returns of individual securities create offsetting positive and negative gains may make it hard to persistently outperform a portfolio without any security selection.

By contrast, even the worst quartile of selection for volatility or correlation still yield positive (albeit below market) Sharpe ratios. This lack of penalty is further enhanced by a payoff for picking the best quartiles of volatility or correlation stocks which show at or above the benefits of being between the second and first quartile of security selection.

The shape of this landscape gives rise to the analysis in the final chapter which will examine the correlation benefit and the effect it has on the panel of allocation methods.

4.4.3.5 Portfolio Combinations Another way to view this isolation by parameter is to examine the scatter plots of all of the resulting portfolios which are formed by the best quartile for the individual parameter.

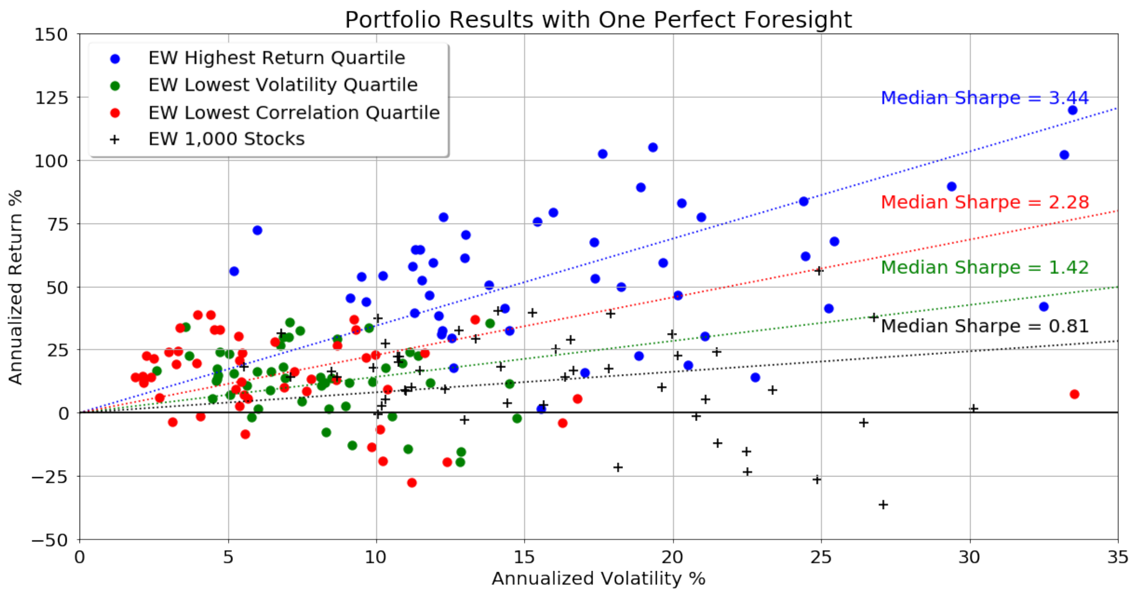

Figure 48: Portfolios Formed With One Perfect Foresight
The figure above displays the annual return vs. risk for portfolios that are formed with one perfect foresight.

The starting point for Figure 48 are the black '+' markers which represent annual portfolios formed by the equal weight of the full 1,000 stocks each year. The blue dots represent portfolios formed by the top quartile of stock returns each year, and have the highest resulting Sharpe ratio. The green dots represent the portfolios formed by the bottom quartile of stock volatility each year. The red dots are the portfolios formed by the bottom quartile of mean pairwise correlations each year.

4.5 Conclusion

It was (Markowitz 1952) who showed the impact of correlations between securities on the risk-return efficiency of a portfolio. Before this time, it was not known how the set of available portfolio could be chosen so as to form a frontier of efficiency. This was a foundational work that clearly showed the impact of these inter-relationships on a total portfolio.

But Markowitz also pointed out the importance of the security selection work, which he did not address in his paper. This chapter of my research adds to the literature by finding a method to measure the relative importance of accurate security selection versus accurate correlation forecasting.

This chapter has demonstrated the dominance of the correlation structure in driving portfolio efficiency. Whereas most practitioner effort is spent on predicting the relative returns of individual stocks, portfolio efficiency gains are significantly more available from harvesting the diversification benefits of a low and stable correlation structure.

It may seem counter-intuitive that by prioritizing security selection in the process of forming portfolios, the diversification benefits can be severely limited. This limitation can not be mitigated through relative performance improvements of security selection.

For both researchers and practitioners, a greater amount of effort should be expended on exploiting correlations as the starting point of portfolio formation. If the generalized hierarchy of correlation can be more accurately forecast, then it can be harvested in the form of higher Sharpe ratios. Portfolios with optimal Sharpe ratios are dominated by harvesting low correlations, not by selecting stocks with high returns.

The next chapter sets out to examine the stability of the hierarchy of correlations and to show the impact of harvesting low pairwise correlations on the panel of allocation methods.

Chapter 5

Correlation Selection

5 Correlation Selection

5.1 Introduction

The previous chapter started from the end point of an optimal Sharpe ratio portfolio given one year of perfect foresight into all stock characteristics (return, volatility, correlation). This approach highlighted that perfectly optimal portfolios are dominated by low correlation stocks, not by high returning stocks. This chapter moves forward with this information to examine the specific effects of selecting stocks based on correlation. This will be termed 'correlation selection'.

The analysis of correlation selection will be taken in three parts, aided by the tool of perfect foresight which is now applied to correlation selection.

- In the first section, perfect correlation foresight is applied to examine the impact on the returns, volatilities and correlations of the available stocks.
- In the second section, a broad range of correlation selection is translated into time series of returns to document improvement or degradation in portfolio Sharpe ratios.
- In the final section, the robustness of the Sharpe ratios differences is tested to show the absolute levels of correlation selection that would be required to change portfolio Sharpe ratios with statistical significance.

The purpose of the analysis is to more clearly understand how stock selection and capital allocation interact. It is not the intent of this research to suggest methods to more accurately select stocks from a variety of methods. But a deeper understanding of the drivers of optimal Sharpe ratio portfolios may help to focus future efforts toward combining robust stock selection with optimal capital allocation methods.

5.2 Refined Method of Selection

In the previous chapter, a simple method of examining stocks by quartiles was shown in Figure 47. In this approach, the full 1,000 stocks in the correlation matrix were sorted by the quartile of their mean pairwise correlation vs. all other securities. While this approach allowed us to see the rough landscape of perfect foresight when applied to correlation, it lacked the level of detail required for a more refined analysis.

While we don't see correlation forecasting much in practice, there is a smaller set of literature which points toward this objective, but with a high degree of mathematical complexity. Given what has been observed with many equity portfolio manager's willingness or ability to use complex mathematical tools, the methods in the literature would certainly be a stretch for them to implement.

For the analysis in this chapter, we need a reliable way to bias the set of available securities up and down just on correlation. While we may have selected from the set of methods seen in the literature, any method chosen would not have reliably created the effect we need in order to study the interaction of correlation directly on portfolio Sharpe ratios. In some years, our chosen method would work well, and in others it would work poorly.

For this reason, I extend the near-perfect foresight method from the prior two chapters to the correlation parameter for our set of stocks. It is acknowledged that this approach is not practical, yet the reasons remain the same as for our near-perfect foresight into returns. Potential diversification is preserved with this approach. In addition, we view this method as a boundary condition such that any other method could not achieve a better result while preserving the diversification benefit. It's purpose is to create a specific set of outcomes so that the relationship between correlation and portfolio Sharpe can be studied with the highest degree of clarity.

For the analysis in this chapter, perfect one year foresight is now applied to correlation, and a broad range of correlation selection is applied. To create the correlation tilts, the process begins with all 1,000 stocks which are ranked by their mean pairwise correlation to all other stocks.

As a first step in creating portfolios tilted toward lower correlating stocks, the stock with the highest ranking of mean pairwise correlation is removed. Conversely, for creating portfolios tilted toward higher correlating stocks, the stock with the lowest ranking of mean pairwise correlations is removed.

The process of removing stocks is conducted one at a time. This method is required because the rankings may change as stocks are removed. A stock may have a high or low mean pairwise correlation, but the impact of its removal on the mean correlation of the other stocks is not known.

Once a single stock has been removed, the rankings are recalculated and the process repeats to remove either high or low ranked stocks to continue the process of tilting the portfolio toward higher or lower correlating stocks. The process is repeated until the portfolio has reached its desired size. In the following sections, both higher and lower correlation tilts are analyzed on portfolio sizes from 1,000 down to 2.

Due to high computational intensity, portfolios which are formed with optimization methods (ER, MV, and MD), are analyzed in increments of 100. The looping process to remove stocks is still completed one stock at a time, but the portfolios are then formed once increments of 100 have been achieved in the selection process.

5.3 Stock Characteristics

As stocks are systematically removed from the beginning set of 1,000 available names, there are observable changes in the characteristics of the remaining stocks. As a first step in the analysis, this section shows the changes in return, volatility, and correlation of the remaining stocks. This is an important first step to understand because in the following sections, portfolios are then constructed from these stocks and it is helpful to understand the inputs to the portfolios. The results at this stage are a cross section of annual observations.

For each of the charts within this section, the perspective of analysis is a single year. Each grey line in the charts represents a single year from the 50 years of history contained within this study. To create each line, a single year of returns is considered, while looping through the selection method described above. This process is followed for portfolios that tilt toward higher and lower correlations. The mean observation across the years is then plotted in a bold color line.

Also, within this section only, as portfolios are not being formed, it was computationally efficient to analyze all sizes of portfolio from 1,000 to 2 in increments of one. This creates a very fine level of detail to analyze changes in stock characteristics.

In a later section, the individual years are compiled to create time series that are examined for characteristics. This will be a switch from cross sectional analysis of years to a time series perspective.

5.3.1 Mean Stock Returns

The first characteristic documented is the change in mean annual returns for the stocks that remain in portfolios with correlation tilts.

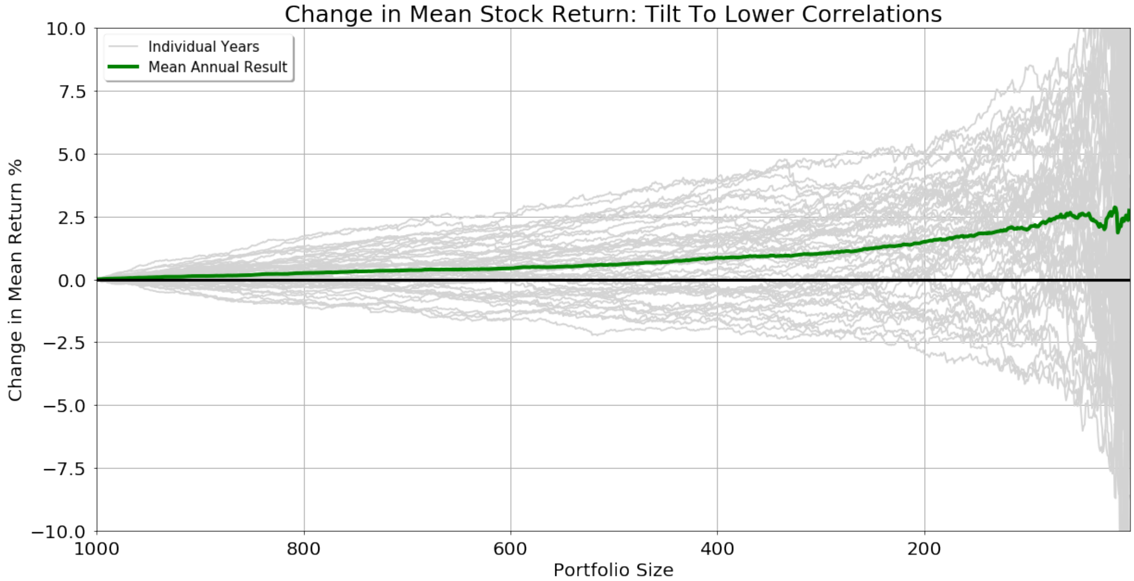

Figure 49: Change in Mean Stock Return - Tilt to Lower Correlation
For each of the 50 years in the study, the chart displays a line showing the change in mean return of the remaining stocks after an increasing number of high correlation stocks are removed with perfect foresight.

While each year of analysis would show a different level of returns, the characteristic to observe is the change in mean returns that occurs from correlation tilting. Figure 49 is designed to show the change in mean stock return on the vertical axis. The horizontal axis shows the size of the resulting portfolio as the tilt is applied.

At the far left of the chart, all of the lines converge to 0% change as there is no correlation tilt applied yet, with a full 1,000 stock portfolio. As the individual year lines proceed toward the right side of the chart, the resulting portfolio size is getting smaller as stocks are removed and the remaining stocks have lower mean pairwise correlation.

A small number of stocks being removed has very little effect on the mean return, but the resulting distribution of changes to returns increases in dispersion as fewer stocks remain in the portfolio. The mean of the observations is also smooth as few stocks are removed but as the remaining portfolio size approaches 100 or smaller, the mean line becomes less stable.

The key observation from this first look is that by removing high correlation stocks, the remaining stocks have higher returns on average. This low correlation anomaly will carry significance as we begin to build portfolios in the later sections of this analysis.

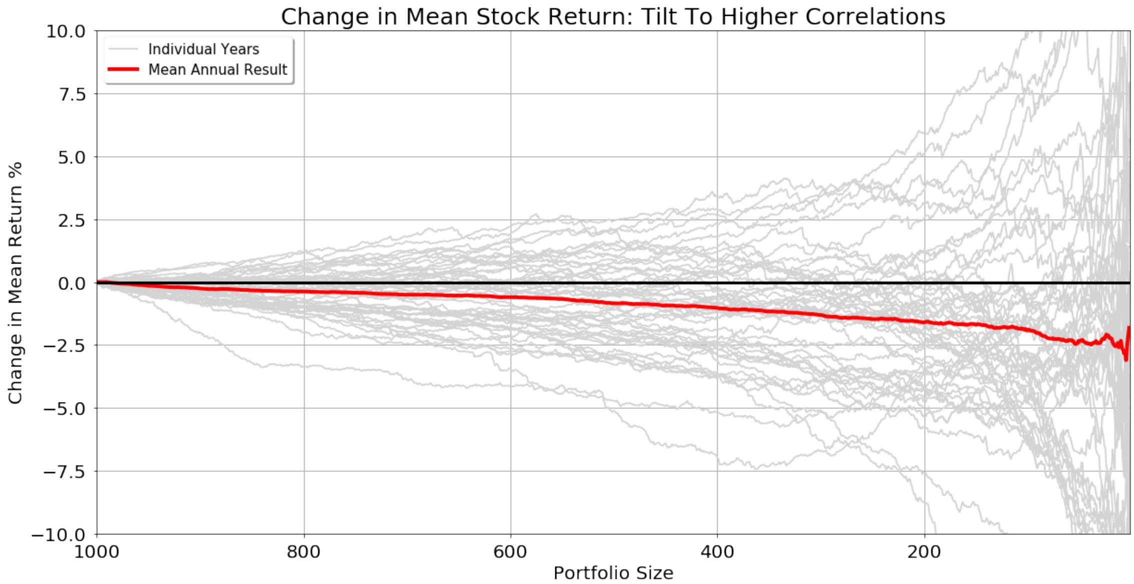

Figure 50: Change in Mean Stock Return - Tilt to Higher Correlation
For each of the 50 years in the study, the chart displays a line showing the change in mean return of the remaining stocks after an increasing number of low correlation stocks are removed with perfect foresight.

While similar to the previous figure, Figure 50 shows the resulting mean return when portfolios are tilted toward higher correlating stocks. As expected, the mean return line falls as the portfolios become smaller. However, the distribution of returns across the years for each size portfolio is dramatically wider and less stable than when tilting toward lower correlations.

In Figure 51, the individual year lines are removed to make clear the spread in returns that occurs from being able to perfectly remove stocks based on high vs. low correlation in the following year. As seen from the two previous figures, not only are the means of

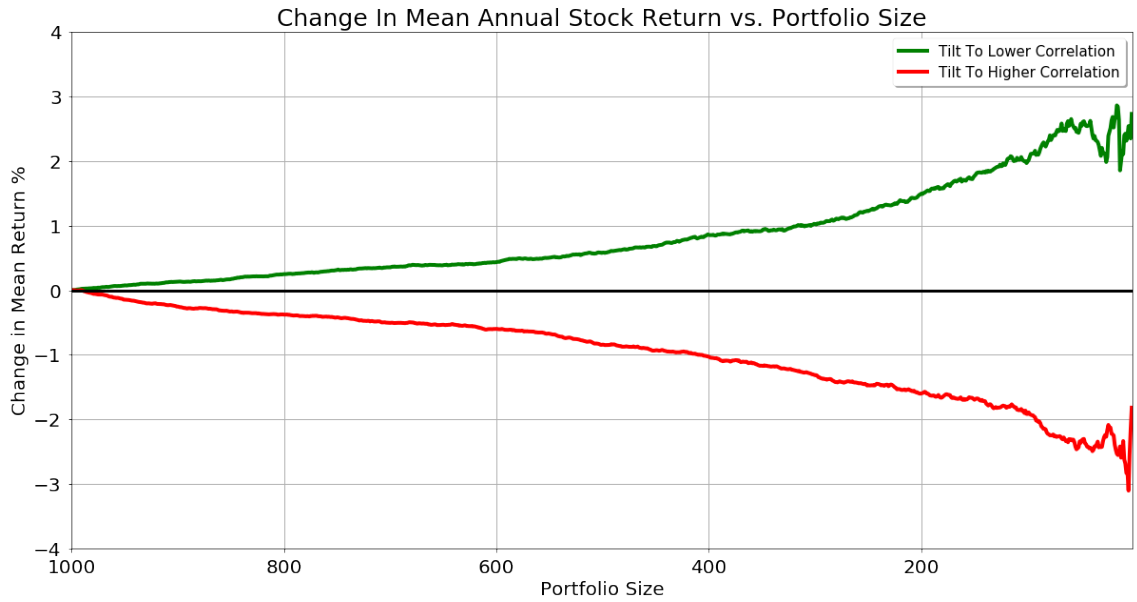

Figure 51: Correlation Selection: Mean Stock Return Comparison
The figure above shows a comparison of the change in mean stock return when tilting toward higher and lower correlation stocks.

the return distributions different, the standard deviation of the distributions is different. Lower correlating stocks have a more narrow dispersion of annual return differences while higher correlating stocks have a wider dispersion of return differences.

If correlation selection is feasible in the absence of perfect foresight, this method would seem to have application to return forecasting for both traditional 'long-only' portfolios as well as 'long-short' portfolios.

In this view of a single characteristic, lower correlating stocks have higher mean returns and higher correlating stocks have lower mean returns. Portfolios built from lower correlating stocks would not only help to increase returns but would also naturally help to decrease portfolio volatility, but only if stock volatility were unchanged by correlation selection. As shown in the next section, this is not an appropriate assumption.

5.3.2 Mean Stock Volatility

Mean stock volatility is the next characteristic to examine in the presence of perfect correlation tilting.

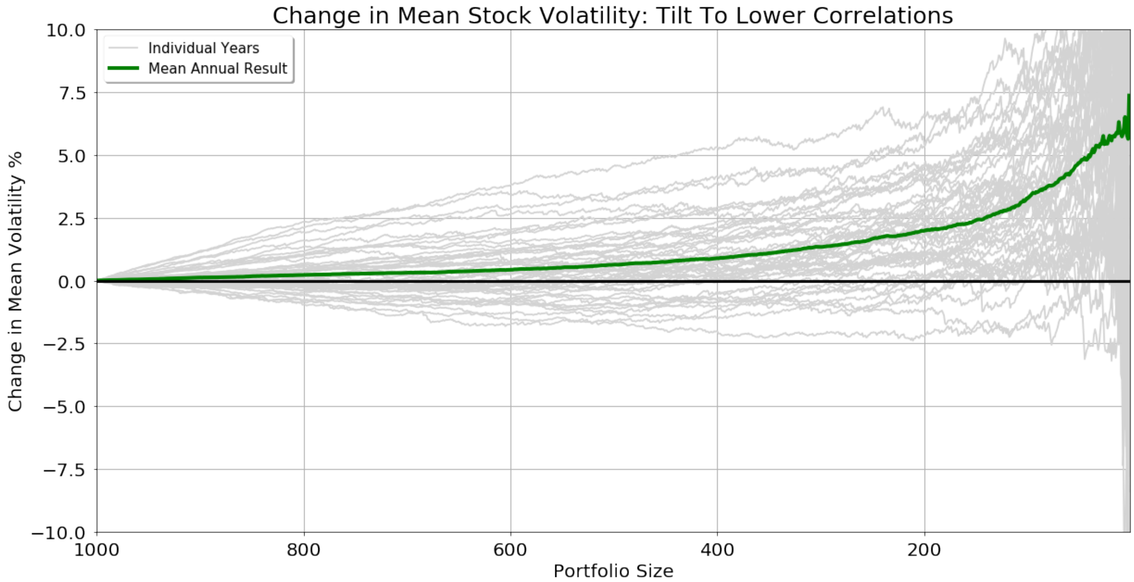

Figure 52: Change in Mean Stock Volatility - Tilt to Lower Correlation
For each of the 50 years in the study, the chart displays a line showing the change in mean annualized volatility of the remaining stocks after an increasing number of high correlation stocks are removed with perfect foresight.

As seen in Figure 52, when tilting stocks toward lower correlations, we observe a relatively smooth increase in the mean volatility of the remaining stocks. The pattern is similar to that of the increase in mean stock returns, but in smaller sizes of portfolios, the effect of increasing volatility is especially present. Also, the distribution of the observations is relatively narrow and only become unstable at portfolio sizes of approximately 100 or less.

By contrast we see in Figure 53 that when tilting toward higher correlating stocks, there is little decline in the mean stock volatility. However, the distribution of the mean stock volatility at each level of portfolio size is dramatically wider.

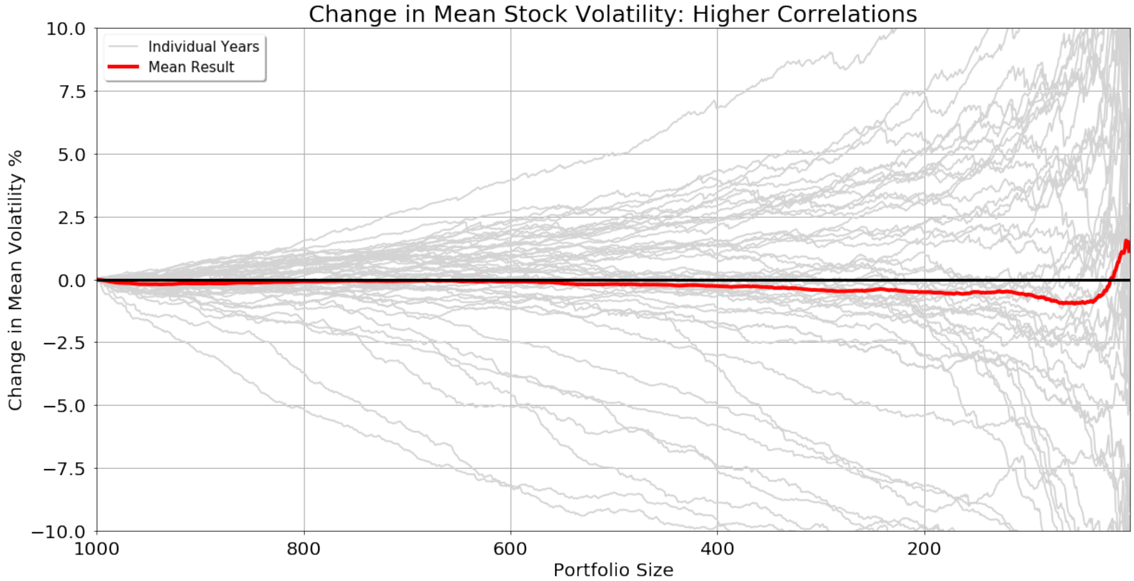

Figure 53: Change in Mean Stock Volatility - Tilt to Higher Correlation
For each of the 50 years in the study, the chart displays a line showing the change in mean annualized volatility of the remaining stocks after an increasing number of low correlation stocks are removed with perfect foresight.

Lower correlation stocks appear to have higher mean volatilities, while higher correlation stocks have a very wide range of volatility and no appreciable difference is observed in the cross section of years when selecting based on correlation alone.

In Section 5.3.1, selecting stocks based on lower correlations improves mean annual returns. However, we can now see that stocks with lower correlations also give rise to higher mean volatilities. These two forces work against each other in the Sharpe ratio of portfolios built with correlation selection. The net impact of these opposing forces within a portfolio will be determined by the correlation between the securities.

Will the lower correlation of the securities be enough to offset the higher mean volatility? Before we answer that question, let's first observe how correlations change when tilting the portfolios.

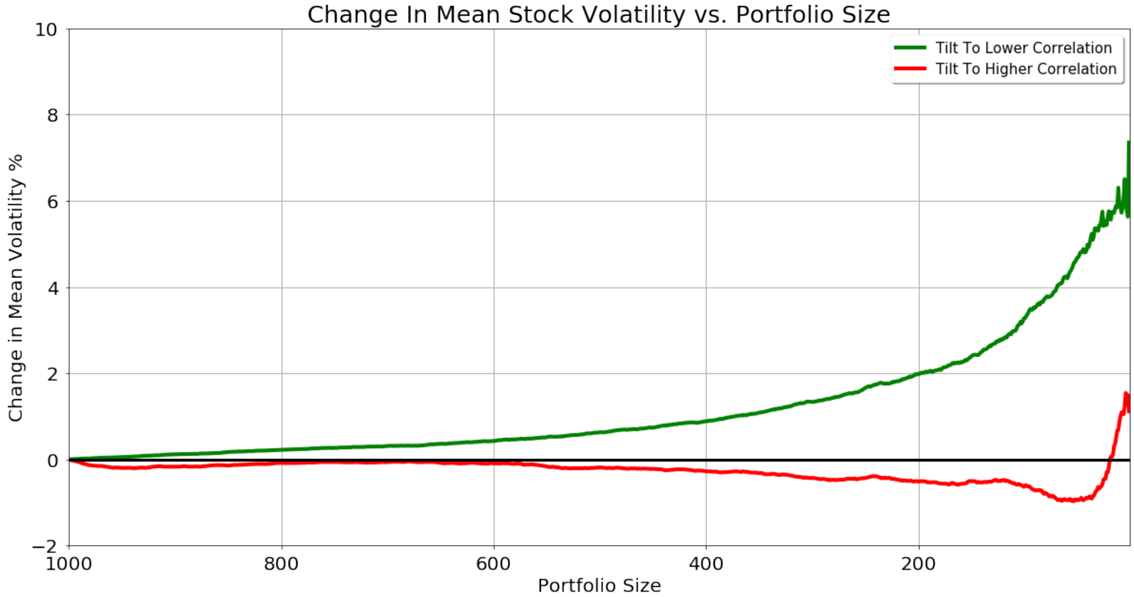

Figure 54: Correlation Selection: Mean Stock Volatility Comparison
The chart shows a comparison of the change in mean stock volatility when tilting toward higher and lower correlation stocks.

5.3.3 Mean Stock Correlations

It is intuitive that as portfolios are tilted toward higher or lower correlations (with perfect foresight), that the resulting mean correlations will shift in the desired direction. For this section however, the method of displaying the change in stock correlations is slightly different. The vertical axis of Figure 55 measures the observed correlations rather than the change measured from the full 1,000 stock portfolio. The starting point of 0.27 should be familiar with observations made in the introduction of Chapter 2 and shown in Figure 2.

Tilting toward smaller portfolios, the observations of the mean correlations for each year proceed in relatively smooth and symmetric fashion toward higher and lower levels. At portfolios of 200 stocks, the spread in mean correlation is approximately 0.30. At the extremes of small portfolios, mean correlations approach 0.8 for higher correlations and

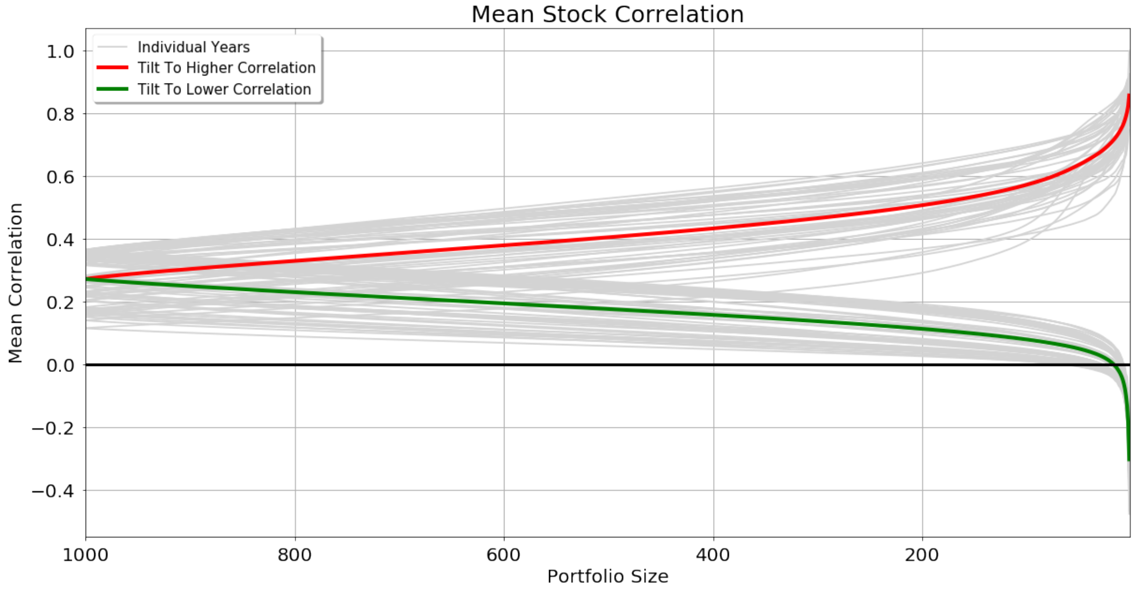

Figure 55: Change in Mean Stock Correlation
The figure shows a comparison of the change in mean pair-wise stock correlation when tilting toward higher and lower correlation stocks.

approach -0.3 for lower correlations. The spread at these extremes is close to symmetric.

5.3.4 Stock Characteristic Summary

This section of the analysis focused on the changes in the mean return, volatility, and correlation of individual securities when perfect foresight is used to select securities based on correlation alone. To conduct the analysis, each of the 50 single years in the study were used, and mean observations were calculated from the cross section of years. The analysis highlighted changes in each characteristic. The next step is to move from mean stock level characteristics to observing the impact on the cross section of annual portfolios.

5.4 Portfolio Characteristics

This next section moves forward from observing the mean characteristics of individual stocks and begins to observe the net impact on portfolios which are constructed from these stocks. As noted in Section 5.3.2, the impact of correlation selection is not easily anticipated in advance. When selecting stocks with lower correlations, returns appeared to increase while volatility also increased. We now combine these stocks together to take the first step in observing the net effect on portfolios.

For this section, a simple equal weighted approach will be used to observe the early differences in portfolio characteristics. In the next section, the full panel of capital allocation methods will be applied in combination with correlation tilting.

5.4.1 Portfolio Return

With an equal weight method applied for this section of analysis, there is no change from the results displayed in Figure 51. The mean observation is created with an equal weight and so in respect to equal weighted portfolio returns, the two concepts are equivalent.

5.4.2 Portfolio Volatility

Changes in portfolio volatility are not quite as obvious. Portfolio volatility is a combination of weights, volatilities, and correlations. In the previous section we saw offsetting results for returns vs. volatilities. Low correlation stocks have higher returns but also have higher volatilities. The correlation tilts will be the deciding factor when observing portfolio volatility.

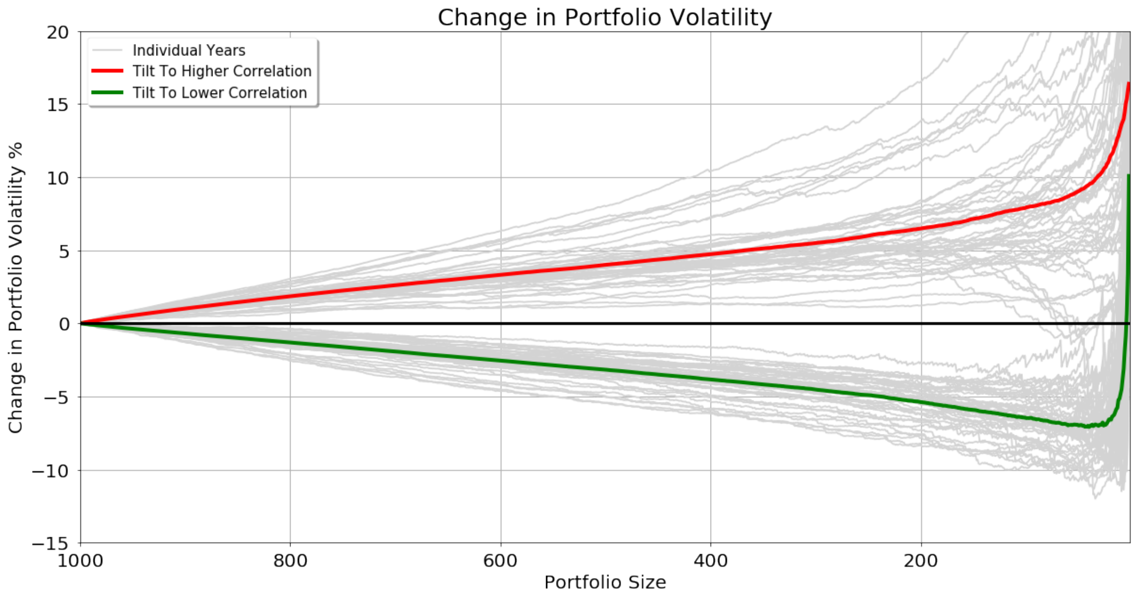

Figure 56: Change in Portfolio Volatility

The figure shows a comparison of the change in portfolio volatility when tilting toward higher and lower correlation stocks.

From Figure 56, it is now apparent to see that the tilts in correlation have a very clear effect on portfolio volatility. Tilts toward higher correlation securities drive portfolio volatility higher despite using stocks with slightly lower mean volatility. Tilts toward lower correlation securities drive portfolio volatility lower despite using stocks with clearly higher volatility.

Changes in portfolio volatility are dominated by changes in correlations. This effect is clear as portfolio volatility moves in the opposite direction of mean stock volatility. At portfolio sizes of 200, the difference in volatility between portfolios tilted toward higher or lower correlation securities can be approximately 10% of annualized volatility. Even in traditional 'long-only' portfolios, correctly forecasting correlations would have a significantly positive impact on portfolio volatility.

5.4.3 Portfolio Sharpe Ratio

The focus of this dissertation work has been to more clearly understand how the elements of security selection and capital allocation interact. A clearer understanding may help to drive work toward improving the efficiency of portfolios, which has been the objective since Markowitz first wrote his seminal paper on portfolio construction.

The next step is to combine the elements observed from perfect correlation selection to display the impact on our objective function, the Sharpe ratio.

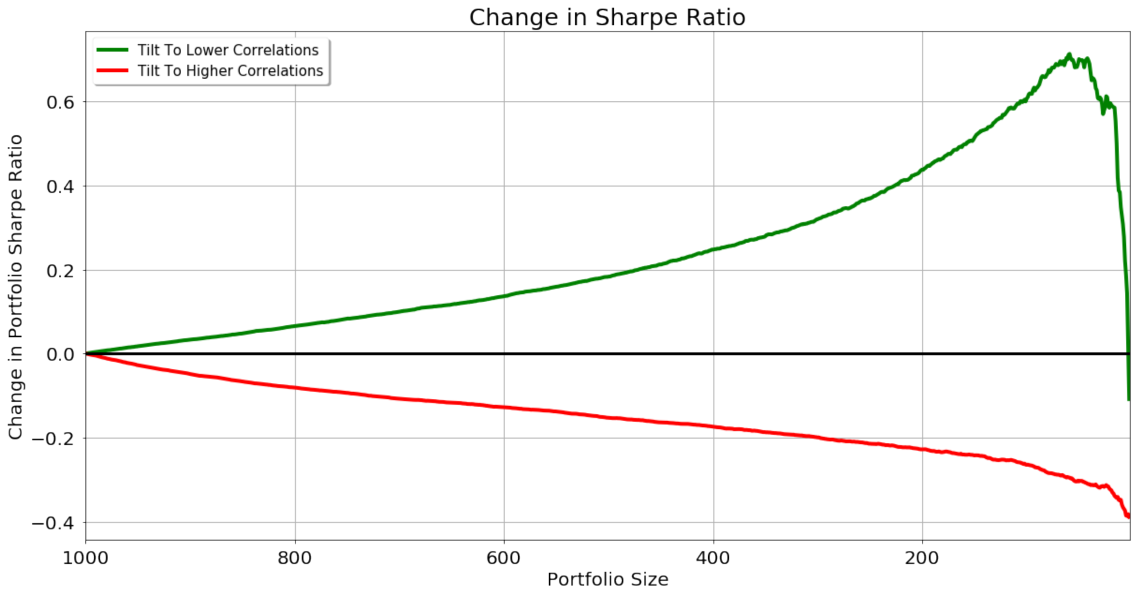

Figure 57: Change in Portfolio Sharpe Ratio
The figure above shows a comparison of the change in portfolio Sharpe ratio when tilting toward higher and lower correlation stocks.

Lower correlating securities have higher observed mean returns and this is a positive impact to the numerator of the Sharpe ratio. Despite having higher mean volatilities, the declining mean pairwise correlations more than offset the higher individual stock volatilities and create portfolios with lower mean volatility. This is a net positive to the denominator of the Sharpe ratio.

The shape of Figure 57 is interesting because it is not symmetrical. This is a reflection of the observation made from the difference in volatilities. But the absolute change in Sharpe ratios is helpful to quantify. For portfolios of 200 securities, 'long-only' portfolios could achieve mean annual improvement in Sharpe ratio of approximately 0.4 with perfect correlation. Conversely, for portfolios of 200 securities with perfectly bad correlation selection would suffer approximately 0.2 decline in mean annual Sharpe ratio.

5.4.4 Portfolio Characteristics Summary

Having moved from observing mean stock characteristics in the cross section of years, this section of analysis combined the stocks into portfolios to observe mean changes in portfolio characteristics. This analysis was also conducted using a cross section of years.

The next step in the research is to make two adaptations. The first change is to move beyond equal weight and apply the full panel of capital allocation methods, while the second change is to pivot from a cross section of annual observations to the creation of time series of returns.

5.5 Capital Allocation Methods

This dissertation has used a panel of capital allocation methods to examine the interaction effect between security selection and capital allocation. The previous sections of analysis focused on a cross section of annual observations and only used the equal weight portfolio as a method for making this observation. In this section of analysis, a pivot is made from cross sectional annual observations to the creation of time series of returns. Within this context, each of the capital allocation methods previously used in this dissertation are applied to quantify the interaction between capital allocation, and securities selected using perfect foresight into correlation.

For each of the charts in this section, a brief explanation of the design is warranted. The horizontal axis displays the time series, beginning from 1968 through 2017 and comprising the 50 years in the study. The vertical axis denotes the growth of $1 over time and is in log scale to reduce the dramatic spread in results which makes the visual observation of volatility more prominent.

Each chart contains four time series. The first is the growth observed when invested in the risk free rate. The second line displays the growth when invested in the allocation method applied to the full 1,000 stocks available. While a wide range of correlation selection has been calculated, the next two lines display the growth when invested in 200 stock portfolios with tilts toward higher and lower correlation securities.

For all methods considered, portfolio sizes between 1,000 and 100 were created in increments of 100. However, displaying the full spectrum of these lines on a single chart becomes quite confusing. Therefore the choice was made to only display two additional lines near the endpoints of the range of analysis. Full details of the returns, volatilities, and Sharpe ratios for each of these portfolios is contained within Figure 75 at the end of this chapter.

5.5.1 Capital Allocation Method - Market Capitalization

This section will highlight and explain the market capitalization method as one of the eight capital allocation methods applied during the analysis. The full explanation for the MC method is given in Section 6.3.1 of this dissertation.

The patterns seen by applying correlation selection tilts with the market capitalization method are similar to those seen with the other seven allocation methods, and so to avoid repetition of observations, the time series charts and summary tables are collected in an appendix at the end of this chapter.

The previous observations made in the cross section of years is evident also in the time series. Portfolios built from lower correlation securities (green line) have higher overall returns in the time series than do portfolios built from higher correlating securities (red line). Visually it is also evident to see the difference in portfolio volatility. The green line is visibly smoother than the red line.

The change in realized volatility is the most striking difference. For example, with the market cap weighted portfolio, simple arithmetic returns are actually higher for portfolios built with a tilt toward higher correlation. However, the significant increase in realized volatility steeply degrades the geometric returns and creates a lower level than portfolios tilted toward lower correlation stocks. This shift is a highlight of the value of lower correlations within a portfolio context.

One exception regarding the pattern of observations is related to the minimum variance and maximum diversification methods. Because these methods are optimizations using the information seen with perfect foresight, there is little to no improvement as portfolios tilt toward lower correlations. But there are more dramatic breakdowns with portfolios that begin to exclude these low correlation securities. The shape of this observation will be made more apparent in the next section.

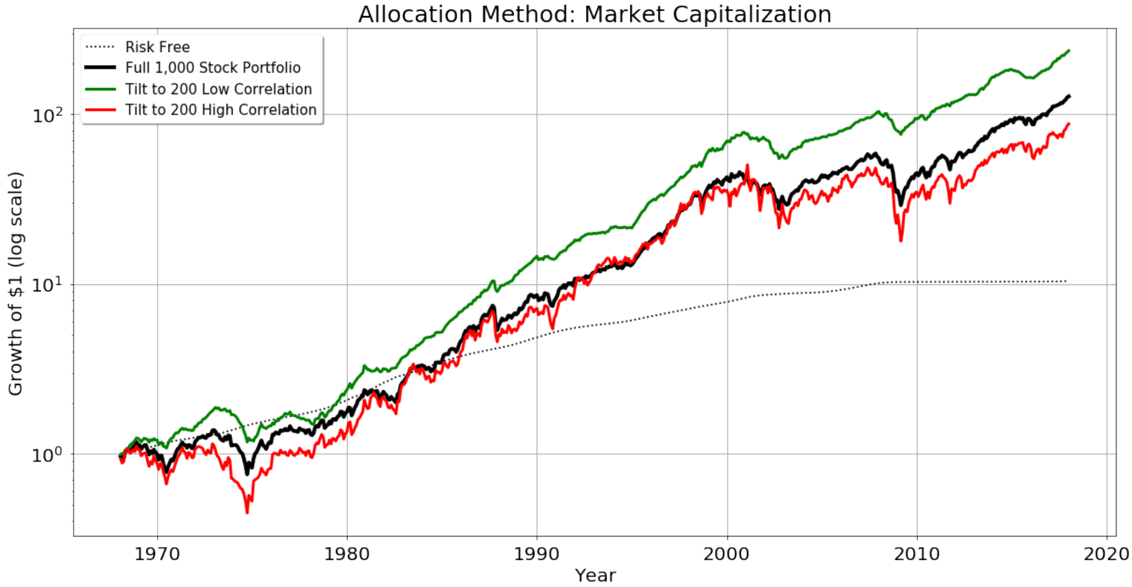

Figure 58: Time Series of Returns: Market Capitalization

The figure above shows a comparison of the time series of returns when tilting toward higher and lower correlation stocks and applying the market capitalization method of capital allocation.

	200 Low	1,000	200 High
Geom Return %	11.55	11.43	9.36
Arith Return %	11.95	10.18	12.77
Volatility %	8.47	15.06	24.98
Sharpe Ratio	0.84	0.44	0.32

Table 9: Time Series Summary - Market Capitalization

The table above summarizes metrics for the allocation method when correlation selection is applied across the 50 year period in the study, 1968 - 2017.

5.6 Sharpe Ratios

The correlation tilts created for this analysis are created each year and on top of these tilts, the panel of capital allocation methods have been applied. Just as we saw in Chapter 3 when security selection was applied with perfect foresight into returns, the Sharpe ratios of the capital allocation methods did not respond symmetrically. Instead, uneven changes were observed across the panel of allocation methods.

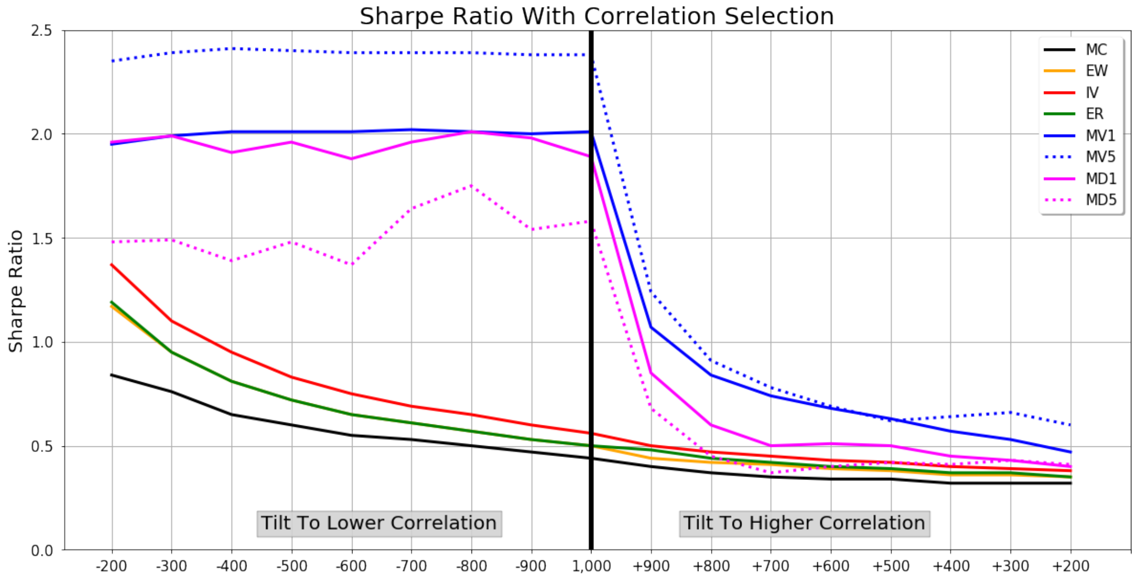

Figure 59: Panel of Allocation Method Sharpe Ratios Across Correlation Selection
The figure above shows a comparison of the Sharpe ratios for each of the capital allocation methods in the study panel and across a range of correlation selection tilts.

Figure 59 shows the Sharpe ratios for each of the capital allocation methods in our study panel and applied across a range of correlation selection tilts. There are several observations that can be made about these results.

First, these methods can be visually sorted into two categories: those with convex shapes (MC, EW, IV, ER) across the full range of tilts, and those with seemingly piece-wise functions (MV, MD). At first this appears quite odd, but the full explanation is a simple matter of the optimization function.

The MV and MD functions optimize for variance or diversification. These capital allocations use a correlation matrix with the benefit of one year perfect foresight. As such, they reach their optimal levels when the full set of securities is available to choose from.

At tilts which exclude higher correlation securities (to the left of the graph), they suffer very little degradation as these aren't the securities that the optimizer is generally choosing. This flat Sharpe ratio pattern continues with little variation all the way down to 200 stock portfolios. This is because these capital allocation methods are constrained by either 1% or 5% maximum positions which equates to 100 or 20 stocks. Each of these levels is below the 200 stock threshold.

Conversely, as the tilt moves toward higher correlation securities (to the right of the graph), the lowest correlation securities are being dropped sequentially and these are exactly the type of securities that these optimized methods choose. As a result, the Sharpe ratio declines rapidly.

By contrast, the other methods are not optimizing an objective function like variance or diversification and do not follow the same pattern of Sharpe ratios across the range of tilts. MC only uses current market capitalization for weights and can improve as lower correlation securities are more available, but does not choose weights based on correlation. EW is a simple equal weight and uses no parameter estimate to choose weights, but does benefit as lower correlation stock are generally more available. Inverse variance does not use correlation for weighting. Only ER is using correlation information, but as variance is a larger component of the covariance matrix, the variance parameter dominates. By definition, ER is placing equal risk into each security and even with a fully known correlation matrix, the method will more closely approximate EW than MV or MD.

A few relative notes are also warranted. For the methods displaying a convex function (MC, EW, IV, ER), the correlation tilts appear to have uniform influence with the exception of MC which displays less benefit at the more extreme tilts toward lower correlation stocks. This observation is in line with a similar observation for MC made in Chapter 3 when we applied perfect foresight to the selection of stocks based upon returns. MC again shows very little responsiveness to an ability to correctly select securities.

With the MV and MD functions, there is an observed difference in the concentration limits. For the MV method, the more concentrated version (5% max position size, dashed blue line) benefits from correctly picking securities based upon correlation. Conversely, for MD, the more concentrated version (5% max position size, pink dashed line) does not do as well as the less concentrated version. Why is this the case?

Again, this is due to the function being maximized during the optimization. Both methods have advance knowledge of correlation and use it in the optimization process. However, MV is minimizing variance, which will pick securities with the lowest combination of volatility and correlation. MD is maximizing the diversification ratio which will tend to pick securities with the best combination of high volatility and the low correlations. So within the low correlation universe, they are picking nearly opposite securities.

Therefore, MD concentrates into higher volatility securities at the cost of Sharpe ratio, while MV concentrates into the lower volatility securities to the benefit of it's Sharpe ratio. This observation also lines up with those made in Chapter 3 regarding MV vs. MD when selecting with perfect foresight into returns.

5.7 Sharpe Ratio Robustness

As a final step in this chapter, it will be helpful to know whether an ability to correctly pick stocks based on correlation would yield robust improvements in Sharpe ratio, or if these increases are just a mirage without justification. If benefits were possible, it would be good to know how accurate the correlation selection would need to be to crate robust improvements in Sharpe ratios. This is similar to the analysis conducted in Chapter 3 for return based stock selection.

Fortunately, we have a useful test developed by (Ledoit and Wolf 2008) and used in Chapter 2 and 3 of this thesis. A description of the test is set in Section 2.4 of this thesis. In brief, the test uses a robust inference method which constructs a studentized time series bootstrap confidence interval for the difference between two Sharpe ratios.

To test our hypothesis that correctly selecting securities based upon correlation alone produces superior Sharpe ratios, we will set the base return stream equal to the returns achieved with the full 1,000 stock portfolio for each method. We then run the MATLAB code provided by the authors to test the difference in Sharpe ratios for each level of tilt and for each allocation method. The test code output provides a p-value which indicates whether the observation is significantly different from the null hypothesis, which in this case is that the true difference in Sharpe ratios is zero.

As an example of the convex functions, we can examine the result of the Sharpe ratio robustness test for the EW allocation method as seen in Figure 60.

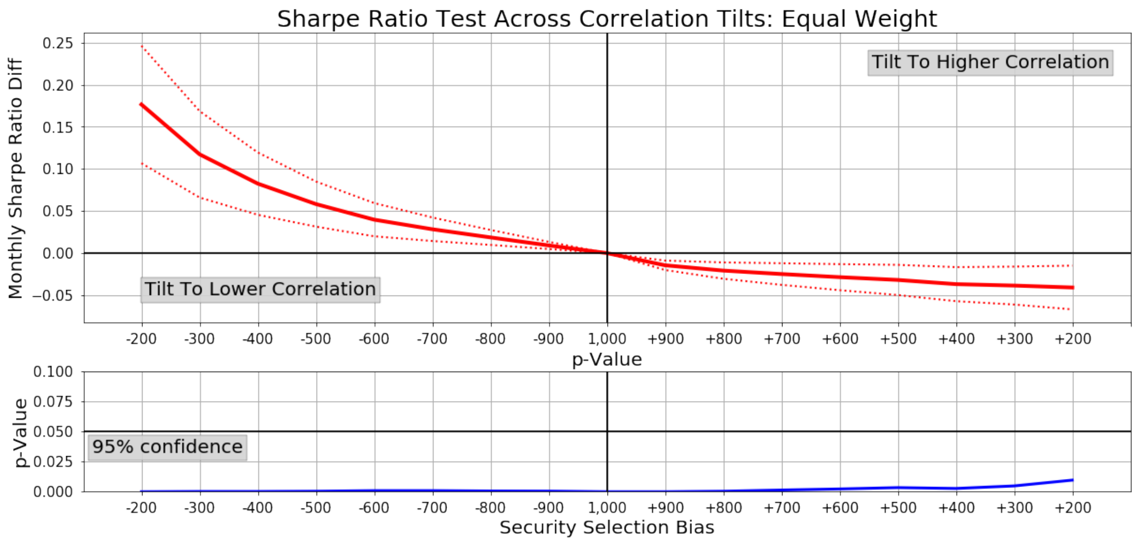

Figure 60: Sharpe Ratio Robustness Test: Equal Weight
The top panel shows the monthly difference in Sharpe ratio ($\hat{\delta}$) across a range of correlation selection tilts with the p-value plotted in the lower panel. P-values below 0.05 are significant at the 95% level.

The top panel shows the change in monthly Sharpe ratio as measured by the robustness test, with a 95% confidence band as calculated by the standard error. In the center of the chart, the lines converge at the full 1,000 stock portfolio and there is no standard error. As the portfolios are tilted to higher and lower correlating securities, the standard error increases and the band around the monthly change in Sharpe ratio ($\hat{\delta}$) widens.

The bottom panel shows the p-value of the Sharpe ratio difference test at each point in the range of correlation tilting. For almost every point in the range of the convex functions, the p-value of these tests sits well below the 0.05 threshold for significance at the 95% level. This indicates that across the broad range, almost every amount of change in Sharpe ratio due to correlation selection is meaningful and statistically robust.

As the Sharpe ratio tests for MC, IV, and ER are all quite similar to the EW test here in Figure 60, the other charts have been included in the appendix of this chapter.

The optimized allocation methods show a different profile in the robustness tests. This difference in profile is echoed in Figure 59. As an example of the difference, the chart for Minimum Variance 1% is shown in Figure 61 while the other variants are included in the appendix.

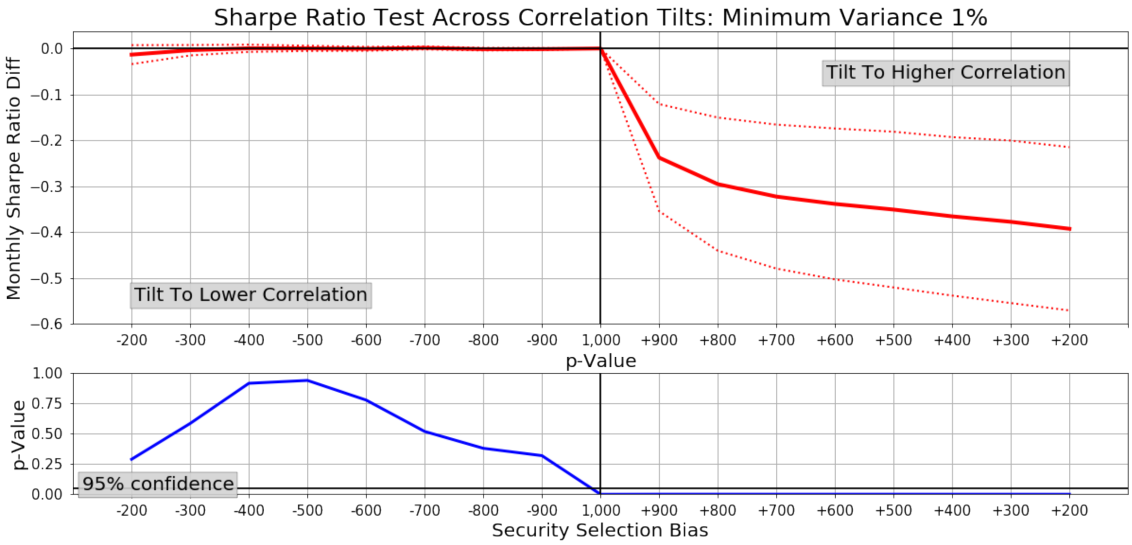

Figure 61: Sharpe Ratio Robustness Test: Minimum Variance 1%
The top panel shows the monthly difference in Sharpe ratio ($\hat{\delta}$) across a range of correlation selection tilts with the p-value plotted in the lower panel. P-values below 0.05 are significant at the 95% level.

With the optimized functions, like Minimum Variance, the left and right side of the panels show very different profiles. On the left side, where the portfolios are tilting toward lower correlation securities and the higher correlating securities are being removed, the difference appears flat. Again, this is due to the fact that the optimized allocation methods all tend toward low correlation securities and have very high concentration limits. The 1% max position size will have around 100 securities while the 5% max position size will have around 20 securities. These are all below the 200 portfolio size limit that is tested in this spectrum.

By removing high correlation securities, nothing meaningful is changing with the securities being selected and so no significant difference shows up in the monthly Sharpe ratio difference ($\hat{\delta}$). Toward the right side of the top panel, the low correlations securities preferred by the optimizer are being systematically removed and the level of Sharpe ratio difference changes very quickly.

In the bottom panel, the p-value displays a very different profile. On the left side with the removal of high correlation securities, despite having very small standard errors, the monthly Sharpe ratio difference is close to zero and even a very small standard error does not allow for a statistically significant difference. The p-values on the left side are quite high and none are significant.

On the right side of the bottom panel, despite having standard errors that are increasing, the monthly Sharpe differences are quite large and still results in statistically robust differences.

As the remainder of the optimized method Sharpe ratio robustness tests are of similar profile, they are all included in the appendix of this chapter.

5.8 Conclusion

The previous chapter extended the literature to allow for a direct comparison between the relative impact on portfolio Sharpe ratios emanating from returns based security selection or correlation based security selection. It was clear from that work that correlations dominated the impact on Sharpe ratios, even to the point of returns based security selection crowding out optimal Sharpe ratios.

This chapter extends the literature by validating these observations with a robustness check. Again, we bring the important work of (Ledoit and Wolf 2008) into the analysis to verify our conclusions.

Building on the awareness from the previous chapter that optimal Sharpe ratio portfolios are dominated by low correlation securities and not high returning securities, this chapter set out to examine the interaction effect of pursuing an optimal process of correlation selection. As in previous chapters, the concept of one year forward perfect foresight was used to examine the interaction effect in the limit. Being able to see the best potential outcome sets an outer boundary condition that limits the expectations for an approach built on the method.

The first step was applying this perfect correlation foresight method to examine the impact on pools of securities. Offsetting insights about how low correlation securities relate to higher returns but also higher volatilities was revealed.

In the next step of the analysis, the range of correlation tilted securities was examined in a portfolio context. This made clear that the low correlations more than offset the higher volatilities as portfolio volatility fell as the correlation tilt increased. The lower portfolio volatility paired with higher returns to create Sharpe ratios that increased. But the increase in Sharpe ratios needed to be tested to check for statistical significance.

To test for statistical significance, time series of returns were created for each allocation method in our panel, and for each level of tilt applied in the range of 200 low correlation to 200 high correlation. The same test method applied in Chapter 2 was useful. In nearly every case, the time series of tilted returns showed a statistically significant difference from the un-tilted (full 1,000 stock) portfolio.

This chapter has demonstrated that securities selected solely on the basis of low correlation can have a positive impact on portfolio Sharpe ratio. This impact can be statistically significant even at very low levels of success. Knowing that optimal Sharpe ratio portfolios are dominated by low correlation securities, while also seeing that even low levels of correlation selection success can drive robust improvements in Sharpe ratio should cause both researchers and practitioners to focus on this approach. Even a moderate inclusion of correlation selection can have a robust impact on portfolio Sharpe ratios.

Chapter 5

Appendix

5.9 Time Series Returns of Capital Allocation Methods with Correlation Selection

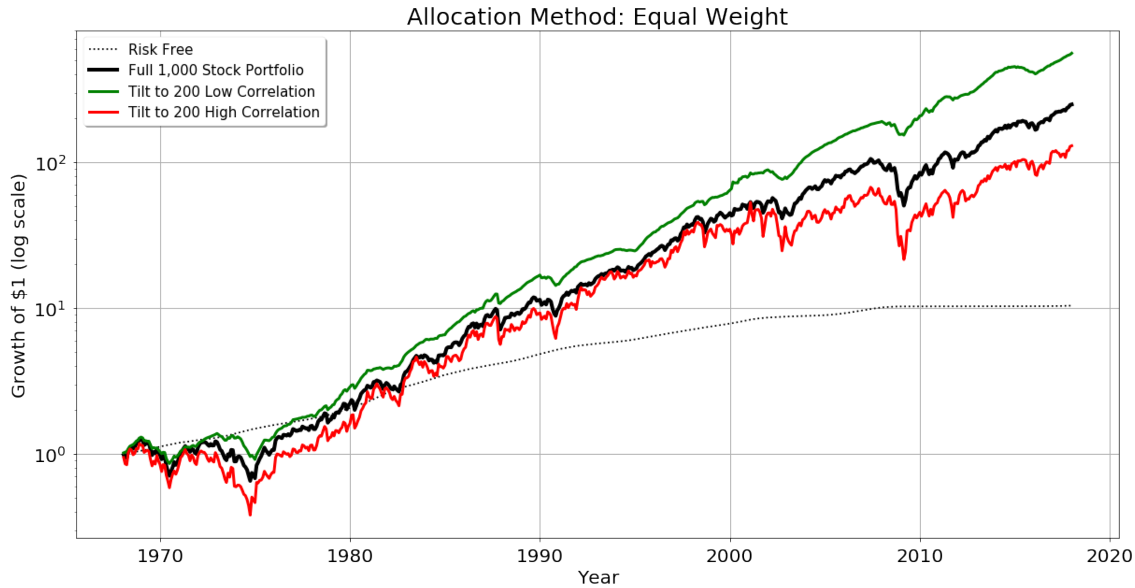

Figure 62: Time Series of Returns: Equal Weight
The figure above shows a comparison of the time series of returns when tilting toward higher and lower correlation stocks and applying the equal weight method of capital allocation.

	200 Low	**1,000**	**200 High**
Geom Return %	13.49	11.67	10.22
Arith Return %	13.82	13.29	14.71
Volatility %	7.68	17.01	28.58
Sharpe Ratio	1.17	0.50	0.35

Table 10: Time Series Summary - Equal Weight

The table above summarizes metrics for the allocation method when correlation selection is applied across the 50 year period in the study, 1968 - 2017.

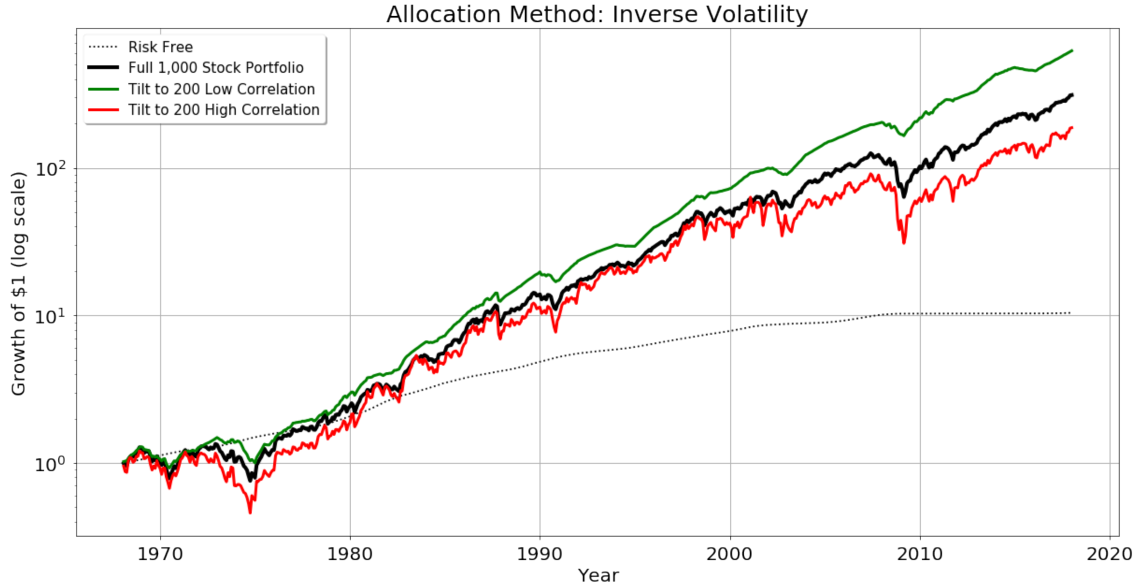

Figure 63: Time Series of Returns: Inverse Volatility
The figure above shows a comparison of the time series of returns when tilting toward higher and lower correlation stocks and applying the inverse volatility method of capital allocation.

	200 Low	**1,000**	**200 High**
Geom Return %	13.73	12.17	11.03
Arith Return %	13.98	13.51	14.98
Volatility %	6.68	15.47	26.70
Sharpe Ratio	1.37	0.56	0.38

Table 11: Time Series Summary - Inverse Volatility

The table above summarizes metrics for the allocation method when correlation selection is applied across the 50 year period in the study, 1968 - 2017.

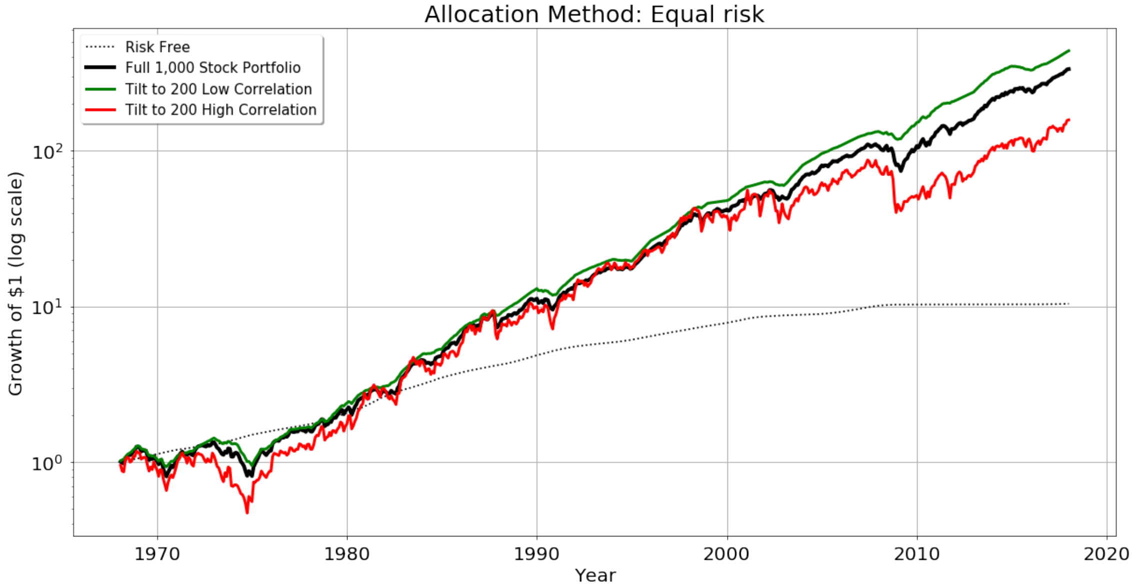

Figure 64: Time Series of Returns: Equal Risk

The figure shows a comparison of the time series of returns when tilting toward higher and lower correlation stocks and applying the equal risk method of capital allocation.

	200 Low	**1,000**	**200 High**
Geom Return %	13.32	11.67	10.37
Arith Return %	13.64	13.29	14.60
Volatility %	7.40	17.01	27.68
Sharpe Ratio	1.19	0.50	0.35

Table 12: Time Series Summary - Equal Risk

The table above summarizes metrics for the allocation method when correlation selection is applied across the 50 year period in the study, 1968 - 2017.

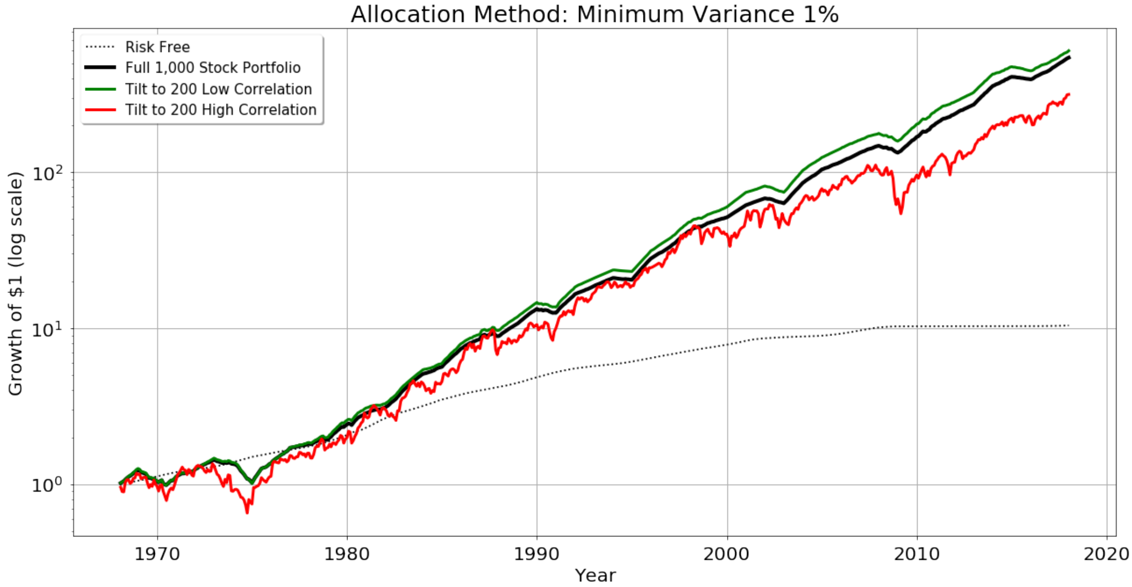

Figure 65: Time Series of Returns: Minimum Variance, 1%
The figure above shows a comparison of the time series of returns when tilting toward higher and lower correlation stocks and applying the equal risk method of capital allocation.

	200 Low	**1,000**	**200 High**
Geom Return %	13.65	13.42	12.19
Arith Return %	13.77	13.52	14.54
Volatility %	4.60	4.34	20.54
Sharpe Ratio	1.95	2.01	0.47

Table 13: Time Series Summary - Minimum Variance, 1%

The table above summarizes metrics for the allocation method when correlation selection is applied across the 50 year period in the study, 1968 - 2017.

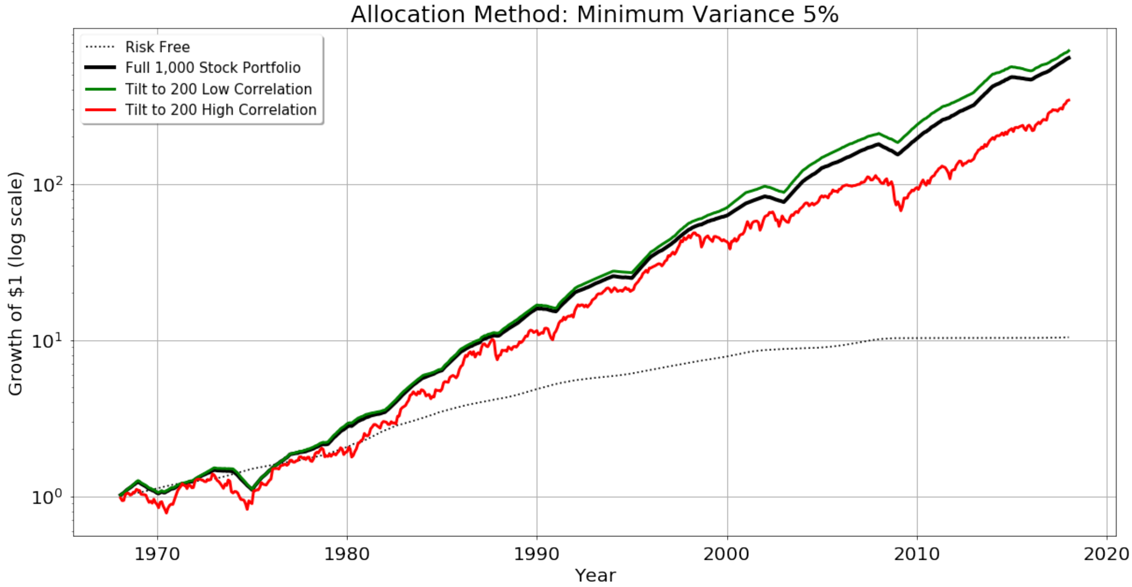

Figure 66: Time Series of Returns: Minimum Variance, 5%
The figure above shows a comparison of the time series of returns when tilting toward higher and lower correlation stocks and applying the equal risk method of capital allocation.

	200 Low	1,000	200 High
Geom Return %	14.03	13.79	12.38
Arith Return %	14.12	13.87	13.59
Volatility %	3.96	3.81	14.72
Sharpe Ratio	2.35	2.38	0.60

Table 14: Time Series Summary - Minimum Variance, 5%

The table above summarizes metrics for the allocation method when correlation selection is applied across the 50 year period in the study, 1968 - 2017.

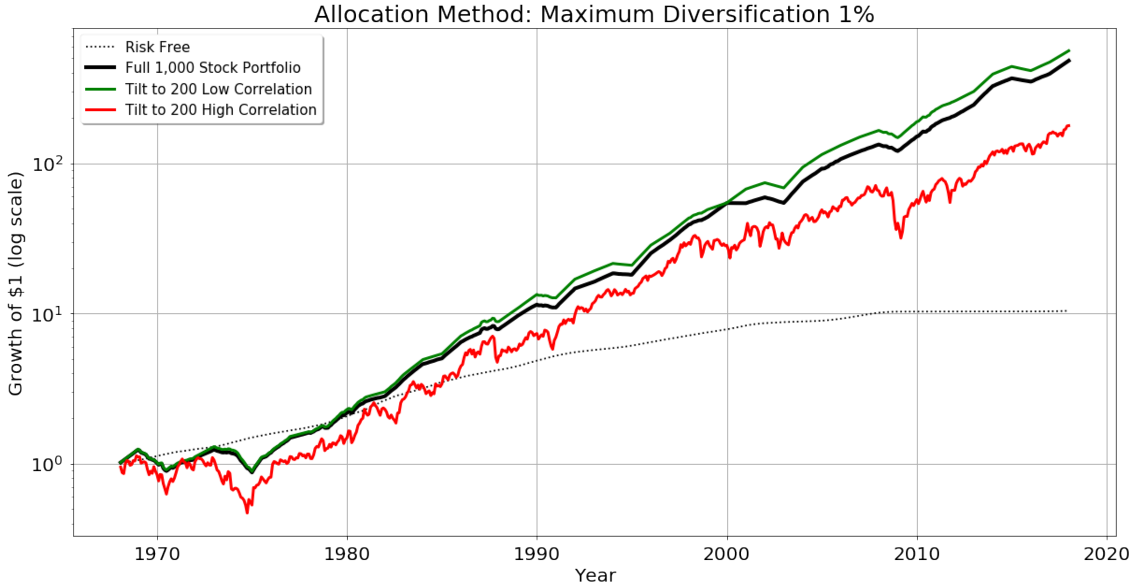

Figure 67: Time Series of Returns: Maximum Diversification, 1%
The figure above shows a comparison of the time series of returns when tilting toward higher and lower correlation stocks and applying the equal risk method of capital allocation.

	200 Low	**1,000**	**200 High**
Geom Return %	13.50	13.15	10.92
Arith Return %	13.61	13.26	13.76
Volatility %	4.50	4.47	22.67
Sharpe Ratio	1.96	1.89	0.40

Table 15: Time Series Summary - Maximum Diversification, 1%

The table above summarizes metrics for the allocation method when correlation selection is applied across the 50 year period in the study, 1968 - 2017.

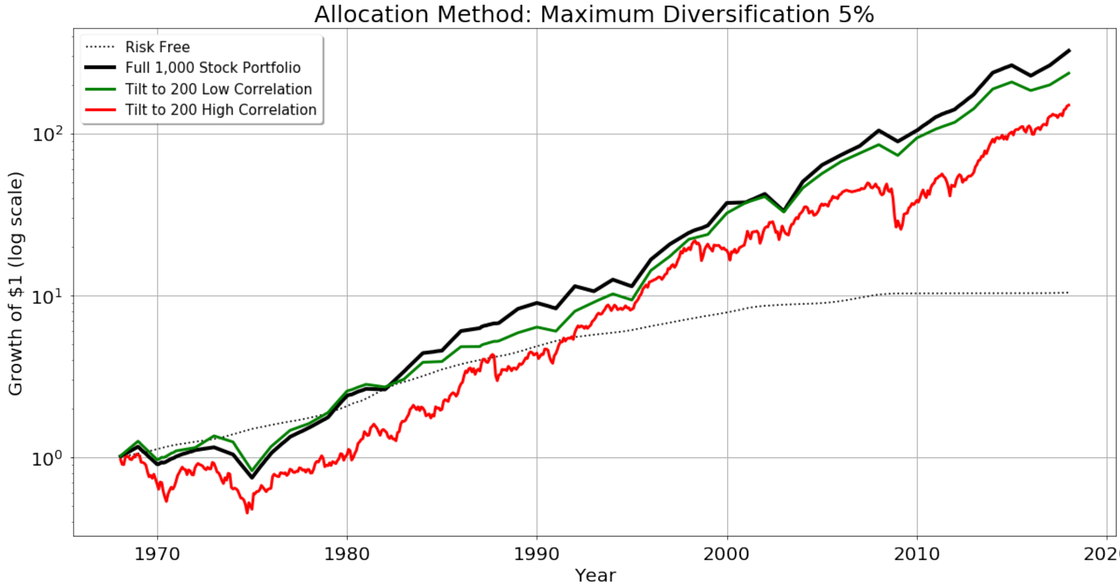

Figure 68: Time Series of Returns: Maximum Diversification, 5%

The figure above shows a comparison of the time series of returns when tilting toward higher and lower correlation stocks and applying the equal risk method of capital allocation.

	200 Low	**1,000**	**200 High**
Geom Return %	11.54	12.26	10.54
Arith Return %	11.66	12.39	12.40
Volatility %	4.63	4.80	18.33
Sharpe Ratio	1.48	1.58	0.41

Table 16: Time Series Summary - Maximum Diversification, 5%

The table above summarizes metrics for the allocation method when correlation selection is applied across the 50 year period in the study, 1968 - 2017.

5.10 Sharpe Ratio Robustness Tests

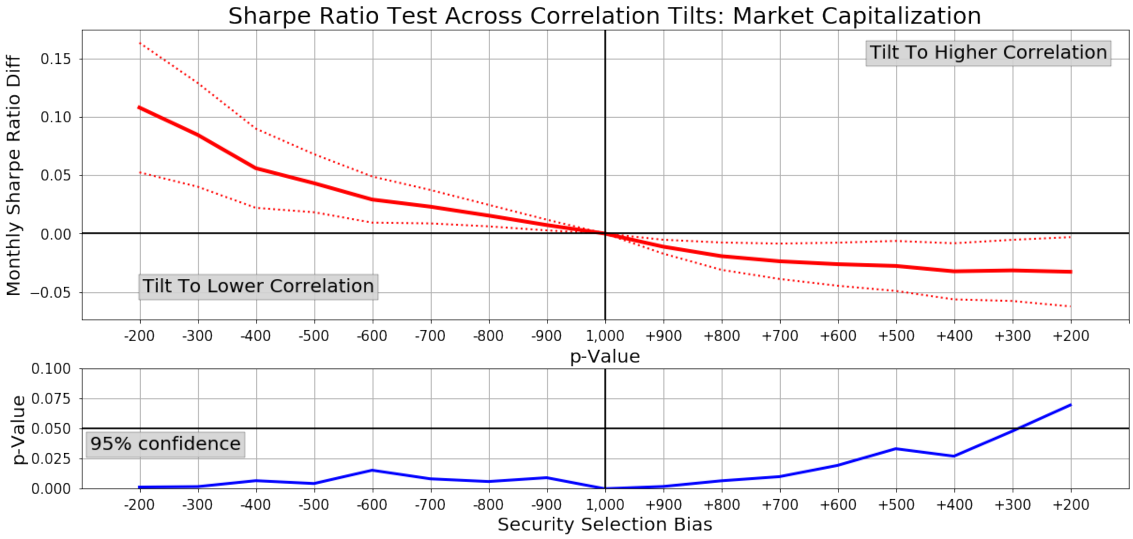

Figure 69: Sharpe Ratio Robustness Test: Market Capitalization

The top panel shows the monthly difference in Sharpe ratio ($\hat{\delta}$)) across a range of correlation selection tilts with the p-value plotted in the lower panel. P-values below 0.05 are significant at the 95% level.

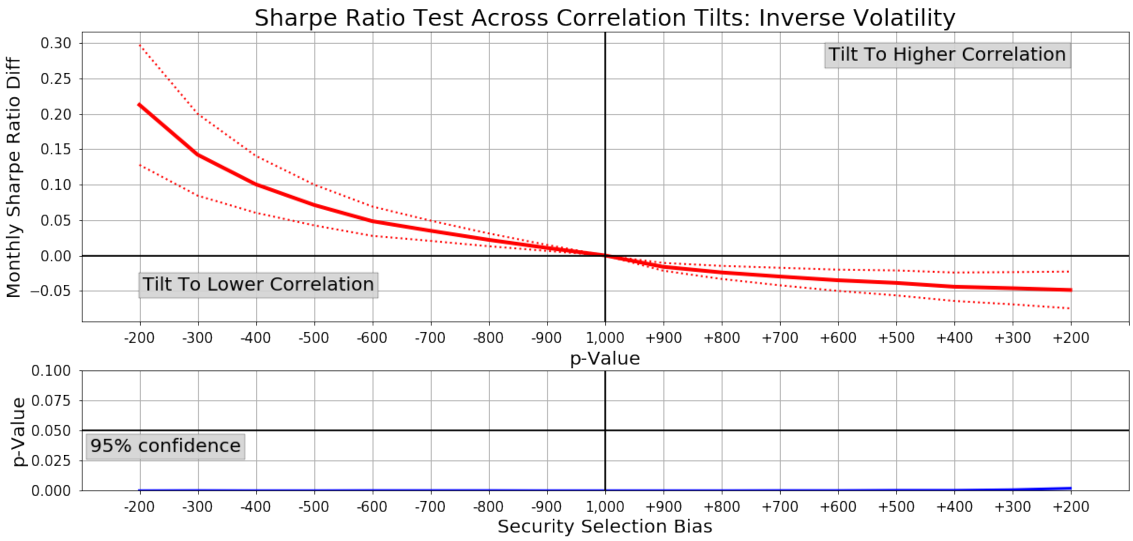

Figure 70: Sharpe Ratio Robustness Test: Inverse Volatility
The top panel shows the monthly difference in Sharpe ratio ($\hat{\delta}$) across a range of correlation selection tilts with the p-value plotted in the lower panel. P-values below 0.05 are significant at the 95% level.

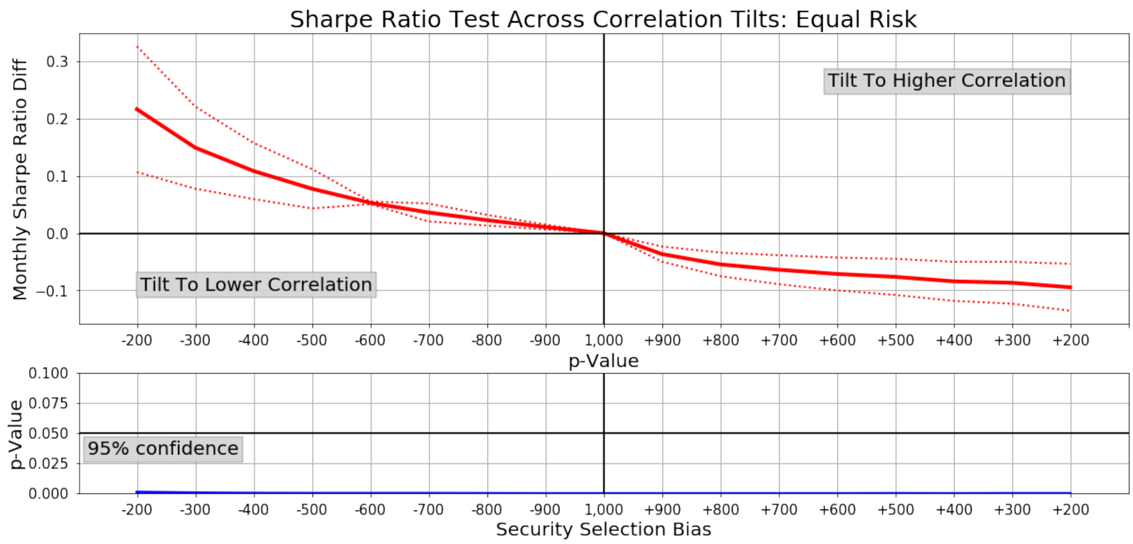

Figure 71: Sharpe Ratio Robustness Test: Equal Risk
The top panel shows the monthly difference in Sharpe ratio ($\hat{\delta}$) across a range of correlation selection tilts with the p-value plotted in the lower panel. P-values below 0.05 are significant at the 95% level.

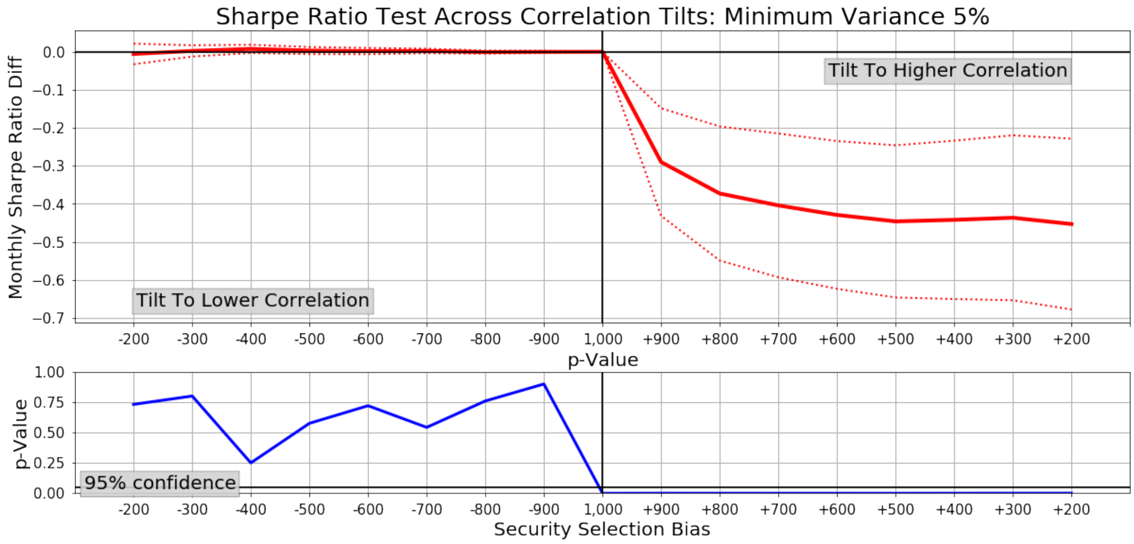

Figure 72: Sharpe Ratio Robustness Test: Minimum Variance 5%

The top panel shows the monthly difference in Sharpe ratio ($\hat{\delta}$) across a range of correlation selection tilts with the p-value plotted in the lower panel. P-values below 0.05 are significant at the 95% level.

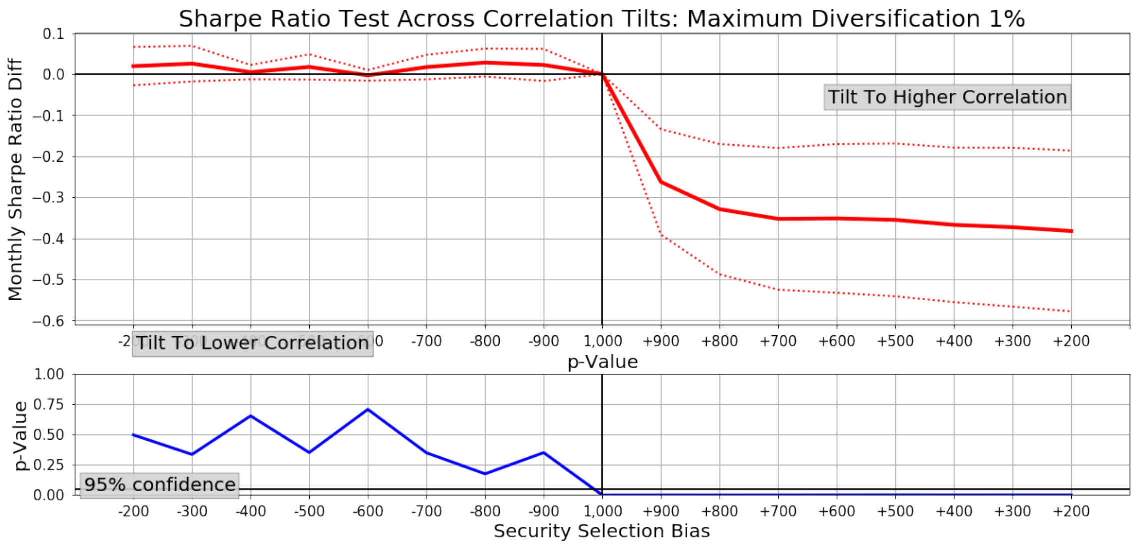

Figure 73: Sharpe Ratio Robustness Test: Maximum Diversification 1%

The top panel shows the monthly difference in Sharpe ratio ($\hat{\delta}$) across a range of correlation selection tilts with the p-value plotted in the lower panel. P-values below 0.05 are significant at the 95% level.

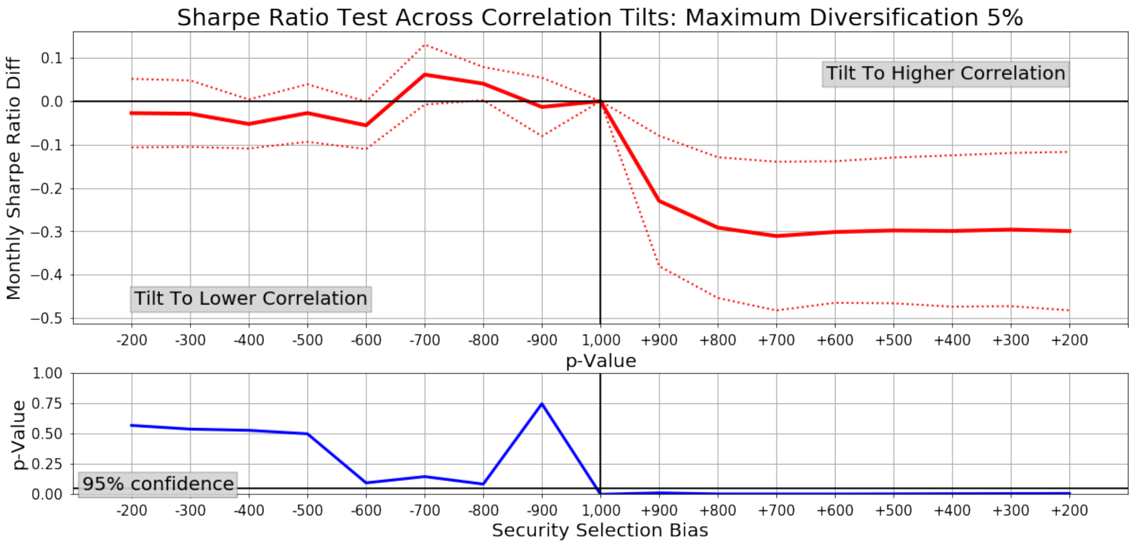

Figure 74: Sharpe Ratio Robustness Test: Maximum Diversification 5%

The top panel shows the monthly difference in Sharpe ratio ($\hat{\delta}$) across a range of correlation selection tilts with the p-value plotted in the lower panel. P-values below 0.05 are significant at the 95% level.

		Tilt to Lower Correlation									Tilt to Higher Correlation							
Portfolio Size		200	300	400	500	600	700	800	900	1000	900	800	700	600	500	400	300	200
MC	Geom Ret	11.55%	11.30%	10.77%	10.74%	10.54%	10.56%	10.47%	10.32%	10.18%	9.85%	9.56%	9.45%	9.41%	9.43%	9.19%	9.38%	9.36%
	Arith Ret	11.95%	11.77%	11.33%	11.41%	11.32%	11.44%	11.45%	11.43%	11.43%	11.33%	11.26%	11.38%	11.56%	11.82%	11.86%	12.42%	12.77%
	Ann Vol %	8.47%	9.20%	10.05%	10.97%	11.86%	12.60%	13.36%	14.18%	15.06%	16.41%	17.57%	18.74%	19.75%	20.82%	22.08%	23.58%	24.98%
	Sharpe	0.84	0.76	0.65	0.60	0.55	0.53	0.50	0.47	0.44	0.40	0.37	0.35	0.34	0.34	0.32	0.32	0.32
	DeltaHat	0.1078	0.0845	0.0559	0.0431	0.0291	0.0230	0.0154	0.0073		-0.0112	-0.0193	-0.0237	-0.0262	-0.0277	-0.0323	-0.0315	-0.0327
	p-value	0.0014	0.0018	0.0067	0.0043	0.0153	0.0082	0.0060	0.0091		0.0019	0.0066	0.0100	0.0193	0.0331	0.0269	0.0476	0.0693
EW	Geom Ret	13.49%	12.87%	12.55%	12.33%	12.11%	12.08%	12.00%	11.86%	11.67%	11.18%	11.02%	10.94%	10.83%	10.72%	10.40%	10.36%	10.22%
	Arith Ret	13.82%	13.33%	13.15%	13.07%	13.00%	13.14%	13.23%	13.28%	13.29%	13.16%	13.31%	13.52%	13.72%	13.92%	13.96%	14.36%	14.71%
	Ann Vol %	7.68%	9.01%	10.27%	11.49%	12.63%	13.76%	14.80%	15.89%	17.01%	18.82%	20.23%	21.51%	22.76%	24.02%	25.38%	26.91%	28.58%
	Sharpe	1.17	0.95	0.81	0.72	0.65	0.61	0.57	0.53	0.50	0.44	0.42	0.41	0.39	0.38	0.36	0.36	0.35
	DeltaHat	0.1765	0.1173	0.0825	0.0582	0.0397	0.0283	0.0186	0.0091		-0.0145	-0.0208	-0.0249	-0.0286	-0.0319	-0.0369	-0.0386	-0.0409
	p-value	0.0000	0.0002	0.0002	0.0004	0.0009	0.0009	0.0006	0.0005		0.0000	0.0004	0.0014	0.0023	0.0034	0.0027	0.0048	0.0096
IV	Geom Ret	13.73%	13.22%	12.93%	12.72%	12.50%	12.52%	12.43%	12.32%	12.17%	11.77%	11.63%	11.52%	11.39%	11.32%	11.07%	11.09%	11.03%
	Arith Ret	13.98%	13.57%	13.40%	13.31%	13.22%	13.37%	13.43%	13.48%	13.51%	13.41%	13.54%	13.71%	13.85%	14.08%	14.15%	14.57%	14.98%
	Ann Vol %	6.68%	7.94%	9.10%	10.22%	11.29%	12.34%	13.36%	14.39%	15.47%	17.14%	18.51%	19.76%	20.98%	22.22%	23.56%	25.06%	26.70%
	Sharpe	1.37	1.10	0.95	0.83	0.75	0.69	0.65	0.60	0.56	0.50	0.47	0.45	0.43	0.42	0.40	0.39	0.38
	DeltaHat	0.2125	0.1421	0.1003	0.0712	0.0483	0.0348	0.0220	0.0106		-0.0161	-0.0240	-0.0298	-0.0351	-0.0388	-0.0442	-0.0463	-0.0488
	p-value	0.0000	0.0001	0.0000	0.0000	0.0001	0.0001	0.0001	0.0000		0.0000	0.0000	0.0001	0.0001	0.0003	0.0003	0.0008	0.0020
ER	Geom Ret	13.32%	12.79%	12.49%	12.30%	12.11%	12.08%	12.00%	11.86%	11.67%	11.65%	11.13%	10.99%	10.89%	10.76%	10.55%	10.55%	10.37%
	Arith Ret	13.64%	13.24%	13.07%	13.04%	13.00%	13.14%	13.23%	13.28%	13.29%	13.23%	13.10%	13.29%	13.50%	13.72%	13.87%	14.31%	14.60%
	Ann Vol %	7.40%	8.87%	10.18%	11.46%	12.63%	13.76%	14.80%	15.89%	17.01%	16.81%	18.75%	20.28%	21.65%	23.06%	24.46%	26.05%	27.68%
	Sharpe	1.19	0.95	0.81	0.72	0.65	0.61	0.57	0.53	0.50	0.47	0.44	0.42	0.40	0.39	0.37	0.37	0.35
	DeltaHat	0.2163	0.1492	0.1084	0.0775	0.0529	0.0360	0.0226	0.0107		-0.0366	-0.0546	-0.0639	-0.0713	-0.0764	-0.0842	-0.0867	-0.0946
	p-value	0.0012	0.0006	0.0003	0.0002	0.0002	0.0002	0.0001	0.0000		0.0000	0.0000	0.0000	0.0000	0.0001	0.0000	0.0001	0.0001
MV1	Geom Ret	13.65%	13.60%	13.59%	13.57%	13.50%	13.52%	13.47%	13.46%	13.42%	11.87%	11.62%	11.67%	11.88%	12.12%	12.15%	12.30%	12.19%
	Arith Ret	13.77%	13.71%	13.69%	13.68%	13.61%	13.63%	13.57%	13.56%	13.52%	12.13%	12.03%	12.24%	12.63%	13.08%	13.14%	13.97%	14.54%
	Ann Vol %	4.60%	4.47%	4.42%	4.41%	4.38%	4.38%	4.37%	4.37%	4.34%	6.86%	8.59%	10.11%	11.59%	13.15%	14.97%	17.28%	20.54%
	Sharpe	1.95	1.99	2.01	2.01	2.01	2.02	2.01	2.00	2.01	1.07	0.84	0.74	0.68	0.63	0.57	0.53	0.47
	DeltaHat	-0.0134	-0.0038	0.0005	0.0003	-0.0007	0.0013	-0.0015	-0.0014		-0.2380	-0.2957	-0.3229	-0.3389	-0.3512	-0.3660	-0.3779	-0.3931
	p-value	0.2876	0.5818	0.9110	0.9346	0.7736	0.5157	0.3776	0.3173		0.0008	0.0008	0.0007	0.0007	0.0007	0.0005	0.0005	0.0003
MV5	Geom Ret	14.03%	13.98%	13.99%	13.94%	13.88%	13.90%	13.85%	13.84%	13.79%	11.41%	10.66%	10.51%	10.37%	10.36%	11.16%	12.26%	12.38%
	Arith Ret	14.12%	14.07%	14.07%	14.03%	13.96%	13.99%	13.93%	13.92%	13.87%	11.58%	10.91%	10.84%	10.79%	10.89%	11.84%	13.16%	13.59%
	Ann Vol %	3.96%	3.88%	3.84%	3.85%	3.83%	3.84%	3.83%	3.83%	3.81%	5.48%	6.73%	7.73%	8.74%	9.84%	11.08%	12.68%	14.72%
	Sharpe	2.35	2.39	2.41	2.40	2.39	2.39	2.38	2.38	2.38	1.24	0.91	0.78	0.69	0.62	0.64	0.66	0.60
	DeltaHat	-0.0057	0.0023	0.0076	0.0032	0.0018	0.0022	-0.0008	-0.0002		-0.2899	-0.3726	-0.4039	-0.4290	-0.4459	-0.4419	-0.4364	-0.4529
	p-value	0.7298	0.7983	0.2495	0.5745	0.7188	0.5411	0.7577	0.8970		0.0007	0.0005	0.0004	0.0003	0.0002	0.0005	0.0009	0.0009
MD1	Geom Ret	13.50%	13.46%	13.16%	13.38%	13.05%	13.41%	13.60%	13.52%	13.15%	10.95%	10.21%	10.06%	10.76%	11.17%	10.98%	11.14%	10.92%
	Arith Ret	13.61%	13.57%	13.27%	13.49%	13.16%	13.52%	13.70%	13.63%	13.26%	11.27%	10.77%	10.86%	11.82%	12.52%	12.69%	13.33%	13.76%
	Ann Vol %	4.50%	4.41%	4.44%	4.44%	4.45%	4.45%	4.44%	4.45%	4.47%	7.59%	10.03%	12.03%	13.79%	15.56%	17.50%	19.83%	22.67%
	Sharpe	1.96	1.99	1.91	1.96	1.88	1.96	2.01	1.98	1.89	0.85	0.60	0.50	0.51	0.50	0.45	0.43	0.40
	DeltaHat	0.0194	0.0255	0.0048	0.0175	-0.0030	0.0171	0.0281	0.0223		-0.2627	-0.3291	-0.3528	-0.3517	-0.3553	-0.3674	-0.3731	-0.3824
	p-value	0.4945	0.3341	0.6507	0.3494	0.7041	0.3474	0.1746	0.3488		0.0008	0.0006	0.0008	0.0014	0.0017	0.0013	0.0015	0.0013
MD5	Geom Ret	11.54%	11.55%	11.39%	11.70%	11.39%	13.37%	13.45%	12.10%	12.26%	8.92%	7.96%	7.76%	8.62%	9.18%	9.55%	10.27%	10.54%
	Arith Ret	11.66%	11.67%	11.52%	11.83%	11.52%	13.49%	13.59%	12.23%	12.39%	9.14%	8.29%	8.23%	9.30%	10.04%	10.65%	11.74%	12.40%
	Ann Vol %	4.63%	4.62%	4.84%	4.76%	4.90%	4.73%	5.03%	4.81%	4.80%	6.34%	7.83%	9.32%	11.13%	12.55%	14.19%	16.27%	18.33%
	Sharpe	1.48	1.49	1.39	1.48	1.37	1.84	1.75	1.54	1.58	0.68	0.45	0.37	0.40	0.42	0.41	0.43	0.41
	DeltaHat	-0.0276	-0.0289	-0.0526	-0.0274	-0.0558	0.0612	0.0403	-0.0134		-0.2300	-0.2915	-0.3110	-0.3016	-0.2980	-0.2993	-0.2960	-0.2995
	p-value	0.5653	0.5351	0.5251	0.4966	0.0938	0.1451	0.0837	0.7432		0.0118	0.0031	0.0028	0.0024	0.0035	0.0047	0.0058	0.0070

Figure 75: Correlation Selection Time Series Summary Statistics

The table above shows the summary statistics for the time series of returns built with the panel of capital allocation methods and a range of portfolio sizes that have been tilted toward higher and lower correlations. DeltaHat is the difference in monthly Sharpe ratios from the bootstrapping test while the p-value is calculated as a two tailed test against the null hypothesis that the true difference in Sharpe ratios is zero.

Chapter 6

Research Summary

6 Research Summary

6.1 Introduction

As an institutional portfolio manager and capital allocator for more than 32 years, I have been perplexed by the process of portfolio formation that I have witnessed in practice. There appeared to be a lack of appreciation for the benefits of systematic capital allocation methods, and for an over-reliance on security selection as a way to improve portfolio returns and Sharpe ratios. A conflict existed in my mind between the literature that I read about the benefits of capital allocation methods and the more commonly employed 'conviction weighting' that I was seeing in practice.

This intellectual conflict spawned a curiosity that led me to begin asking a series of questions which developed during the course of this research project. Each of the question in turn was addressed as a chapter of this thesis, and below I summarize these questions with my observations from the research work.

But good questions are more valuable than good answers. Sometimes that is because they uncover things not previously discovered, but in the way that I think about the value of a good question, the greatest value comes from the additional questions it uncovers. My work is not a conclusion for these topics but a slight extension of the existing understanding. My hope is that these questions and observations give rise to more good questions that can drive us all forward in the practice of portfolio formation and management.

6.2 Hierarchy of Sharpe Ratios

Much of the existing literature on capital allocation methods point toward improvements in portfolio efficiency. However, most active US equity managers that I have encountered do not use these systematic capital allocation methods but instead rely on their conviction which is formed during the security selection process.

Security data over the past 50 years shows that stock returns were widely dispersed giving rise to potential improvements in portfolios through returns based security selection. But the mean annual pairwise correlations are also quite low and stable giving the potential to create substantial improvements in portfolio Sharpe ratios.

My expectation was that these two stages of portfolio formation must both be important and connected somehow. But in order to first understand whether the purported value of systematic allocation was true, I needed to test the robustness of the hierarchy of Sharpe ratios across a commercially available panel of systematic capital allocation methods.

While holding security selection fixed, and using the full 1,000 largest capitalization stocks available each year, I establish an initial hierarchy of Sharpe ratios across the methods. As seen in previous literature, the alternative allocation methods all appear to be better than the market cap method, but were these robust? An application of Ledoit's studentized bootstrap robustness test, the initial conclusions begin to fade a bit. Only about half of the initial observations are robust.

But this was just the first step, and only established the baseline. The question remained about the connection between security selection and capital allocation method.

6.3 Contours of Sharpe Ratios

The existing literature was focused on either security selection techniques or capital allocation methods. My question focused on the interaction effect between these two stages of portfolio formation and this type of testing had not been seen in the literature. This next step forward created a bridge that spanned these two strains of the literature and allow me to examine the dependency of capital allocation on the success or failure of security selection.

To move into the next section of research, I needed to vary the set of securities available to the allocation methods each year. While there is a very wide variety of methods available for security selection, the required method would need to create a reliable bias each year in the return profile of the available stocks. Fundamental and quantitative methods were examined. While discretionary approaches are often observed in practice, they don't make for a robust method of examination. A systematic method was needed in order to create the desired effect in a reliable manner.

A simple method of near-perfect foresight was crafted to examine this question. While it is not achievable in practice, it is systematic and reliably creates the effect needed to study the interaction between selection and allocation. In addition, the method created a boundary condition where no other methods could be created to achieve the same reliability while maintaining the potential for diversification.

This simple security selection method was applied across a range of positive and negative security return bias, giving the opportunity to see the landscape of selection impact on portfolio Sharpe ratios.

The key insight from this part of the research was that the hierarchy of Sharpe ratios was not invariant. Picking stocks well or poorly has an uneven impact on the Sharpe ratios of the various capital allocation methods. The Sharpe ratios of some allocation methods responded in a non-linear way, and other responded in a linear way, but the rate of change varied across the strategies.

This non-invariance across the range of security selection was then tested with the same robustness methods used in the previous chapter. In most cases, the initial conclusions about the hierarchy of Sharpe ratios changed at some point in the range of security selection bias. In a few cases, the conclusions completely inverted toward the ends of the spectrum on security selection bias.

It was shown from this work, that the optimal choice of a systematic capital allocation method is dependent on the ability to successfully or poorly pick securities. Some allocation methods help to amplify good security selection skills while others help to moderate the impact of poor security selection. Knowing this interaction effect is important to creating a better process for portfolio formation and management.

6.4 Low Correlations Dominate

With a more robust understanding for the interaction effect between selection and allocation, my next question was whether returns or correlation were more important toward creating better Sharpe ratios. It appears that active US equity managers spend a large amount of time focused on trying to select stocks that they think will outperform, but not nearly as much time on how to allocate their capital across their selections.

Thinking of the Markowitz efficient frontier, picking stocks with higher returns lifts the entire frontier upward in the risk-return space. Picking stocks with the lowest correlations has the effect of creating a more curved frontier. While this impact may be intuited, there is little way to intuit the interaction effect. Does picking stocks with higher returns create a set with higher or lower correlations? Does picking a set with lower correlations generally increase or decrease median returns?

To study this interaction effect and answer the question of which element would have a greater impact on lifting portfolio Sharpe ratios, I extended the logic from the previous near-perfect foresight method and expanded it to this problem by now looking forward and entire year on all characteristics of the stocks; return, volatility, and correlation.

With a full year look ahead on all elements, I then created Optimal Sharpe Ratio (OSR) portfolios which were free to pick any combination of stocks to have the highest Sharpe ratios. Again, this methodology is not practical for implementation, but it is highly efficient to study the question in focus.

As a comparison, I also designed portfolios which would pick the 20 best performing stocks each year and equal weight their holdings. This concentration of holdings was selected as it was both inline with the most concentrated methods in my panel of allocation methods (MV5% and MD5%), as well as being inline with concentrated portfolios that are commercially available to investors.

If picking the best performing stocks was the preferred path to higher Sharpe ratios, these concentrated portfolio should be near equal with the OSR portfolios. But the risk adjusted performance comparison wasn't remotely close. Median annual Sharpe ratios from OSR portfolios were nearly 23x the Sharpe ratios seen from the concentrated portfolios built from the best performing stocks.

This was an interesting set of data points which indicate that picking top performing stocks wasn't the best way to go when trying to maximize Sharpe ratio. But we needed to understand what kinds of stocks the OSR portfolios were picking. For this question, the analysis moved to a time series to see each year how the selection sets compared.

With regard to return, the concentrated 'best returning stocks' portfolios had vastly superior returns, as expected. The OSR stock returns were a little bit better than the full 1,000 stock universe each year. For volatility, high returning stocks had a little more volatility on average, but the OSR stock returns were virtually the same as the full 1,000 starting stock universe. However, while the concentrated 'best returning stocks' had virtually identical mean pair-wise correlation to the full 1,000 starting stock universe, it was the OSR stocks that displayed very low mean pair-wise correlations.

The 'best returning stocks' portfolios contained 20 holdings but had been constrained to be equal weighted. As an attempt to help this portfolio even further, once the best returning stocks had been selected, I then let the optimizer try to create the highest Sharpe ratio by letting it look ahead a year to see the volatilities and correlations. This would free the portfolio to maximize this information as a second step in the process. While Sharpe ratios increased, they could not get anywhere close to the Sharpe ratios from the portfolios that did not engage in security return selection as an initial step.

This is a very important observation as it indicates that even with exceptionally good security selection based on future returns, this crowds out the potential for maximum Sharpe ratio. But this is exactly what most equity managers are doing. They are trying to pick the best returning stocks first and then allocate capital second. This approach has the effect of limiting their ability to maximize Sharpe ratios.

This analysis points toward the relative importance of correlation versus returns in building portfolios with the highest Sharpe ratios. But the stocks selected by the OSR methodology were the best combination of all three. Following just the OSR method doesn't isolate correlation alone, it only indicates that this is the predominant characteristic of the stocks chosen for the portfolios with the highest Sharpe ratio potential.

As a setup for the final stage of analysis, stocks were sorted into quartiles based first on returns, then volatility, then correlation. Each quartile was compared for the resulting

Sharpe ratio.

Stocks with the highest returns in general created the highest Sharpe ratios, but the curve of reward vs penalty was very steep. Good picks here are easily offset with bad picks. With the volatility sorting, stocks with lower volatility created portfolios with higher Sharpe ratios. This is a portfolio extension of the low volatility anomaly that had been previously documented in the literature. While low volatility stock created portfolios with higher Sharpe ratios, even high volatility stocks created positive Sharpe ratio portfolios, and the curve of reward vs risk was not nearly as steep as that seen with stocks sorted by returns only.

Interestingly, stocks sorted by correlation quartile also displayed the same qualities as the low volatility anomaly. However, we see in the shape of the reward vs risk curve an even steeper impact for the quartile with the lowest correlating stocks.

This step in the research process more clearly highlighted the relative importance of low correlations rather than high returning stocks. But the simplistic approach of using quartiles pointed me toward the final section of research where the analysis would be refined and become more complete.

6.5 Correlation Selection

The final stage of analysis needed to isolate the impact of security selection based on correlation, rather than security returns. We see from the previous chapter that the best Sharpe portfolios are dominated by low correlations, but if that type of selection were possible, what impact would it have on returns and volatilities of the stocks and on portfolios through time.

The first type of analysis conducted was examining the effect of near-perfect correlation selection on the cross section of stocks. Using one year look ahead bias to correlation only, we saw that lower correlation stocks tended toward higher returns, while higher correlation stocks tended toward lower returns. Taken as a single observation, this would seem to be quite beneficial with higher returns and lower individual stock correlations which should lead to lower portfolio volatility.

However, using the same near-perfect correlation selection, we saw that low correlation stocks tended toward higher volatility, while higher correlation stocks had a very slightly lower volatility. This complicated the first observation with regard to returns.

The objective was not to study individual stock characteristics but to observe the impacts on portfolios. As we tended toward lower correlation securities, we saw returns increasing and volatility increasing. Combining these securities into portfolios would allow us to see if the low correlations helped to net out the impact of higher volatilities. Indeed, a combination of lower correlation stocks did create portfolios with lower volatility, rather than higher volatility that we had seen in the cross section of stocks. This shows that correlations are playing a larger role in portfolio Sharpe ratios.

The last part of the cross sectional analysis was to examine the Sharpe ratios and here we see a large increase in Sharpe ratios for portfolios comprised of lower correlation securities.

The next section of analysis on correlation selection was to transform the cross section into time series analysis. For each of the capital allocation methods in the panel, they were formed using three sets of stocks. First was the full 1,000 stocks, then the 200 lowest pairwise correlation stocks, and finally the 200 highest pairwise correlation stocks.

The observations in the cross section transferred well to the time series. Portfolios built with lower correlation stocks displayed higher returns, lower volatility, and higher Sharpe ratios. In addition, the compounding effect of a lower volatility return stream was evident. This was not observable from the cross section.

Finally, the impact of correlation selection was studied across the panel of allocation methods. Just as was observed with the selection of securities based upon returns, the Sharpe ratios of portfolios built from lower correlation securities had uneven reactions.

While it seems that correlation selection could be quite helpful to portfolio Sharpe ratios, it was not certain how good you would have to be at this skill in order to make a robust difference in the resulting Sharpe ratio. Fortunately, the Ledoit robustness test could be used here as well. The net result of all robustness tests shows that correlation selection at almost any level has a positive and robust impact on portfolio Sharpe ratio. This is in contrast to what was observed with security return forecasting. With return forecasting, you had to exhibit substantial skill to make a robust positive impact on Sharpe ratios.

6.6 Summary

My questions arose because of a conflict that I observed between the academic literature and the practice of portfolio formation. These questions led me down a path to examine:

- Were the systematic allocation methods robustly better?
- How would security selection impact the choice of an allocation method?
- What stock level characteristics would dominate the best Sharpe ratio portfolios?
- How much correlation selection skill is needed to robustly improve Sharpe ratios?

Each of these questions led to observations that used existing approaches to build upon previous literature and expand the knowledge about the dynamics of portfolio formation. It is my hope that this work will help practitioners to close the conflict between academic observation and practice, leading to better portfolio outcomes for investors. Additionally, I hope that this work aids researchers who will have new and better questions to build upon this work.

6.7 Future Research

From this thesis, crafty researchers will certainly be able to ask even better questions and follow the path forward. While many questions may arise, a logical extension of this research would be to combine the observations herein with the literature on correlation forecasting and the stability of correlation ranks. From the final part of analysis in Chapter 5, we can see that any ability to improve security sets based on correlation can bring a robust improvement in Sharpe ratios. Finding a way to bring this observation into practice would certainly be a benefit to investment managers and their investors.

Appendix 1

Data & Methods

7 Appendix 1

7.1 Introduction

Across each of the chapters there are common sets of data and methodologies which must be fully described in order to proceed. To make each of the main chapters more focused, all of these elements are described in this appendix. In section 7.2 the sources of the input data are described. Section 7.3 describes in detail each of the allocation methods and how they are calculated. Finally, section 7.5 compares the return, statistical characteristics, and factor loadings of each of the allocation methods.

7.2 Data

7.2.1 Sources

Multiple data sources are needed to construct the portfolios using the panel of allocation methods. Equity data (returns, market capitalization, prices, and share classifications) are collected from the CRSP/Compustat merged database as made available from the Wharton Research Data Service (WRDS) [11] the website of Kenneth French at Dartmouth University[12] is the source for 1 month US Treasury Bill rate (as provided by Ibbotson & Associates, Inc.) and for factor returns for market, size, and value.

[11] `https://wrds-www.wharton.upenn.edu`

[12] `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/f-f_factors.html`

7.2.2 Time Frame

To calculate the weights for the allocation methods which require an optimization routine, 5 years of historical monthly returns are used to estimate the covariance matrix. Figure 76 displays the number of stocks available in the CRSP database since 1951. This study uses the 1,000 largest capitalization stocks available each year. The starting date of 1968 was chosen which allowed for at least 2,000 stocks with 5 years of history. Even in the early years of the study, this starting date provides enough stocks to screen at most the top half as "large cap".

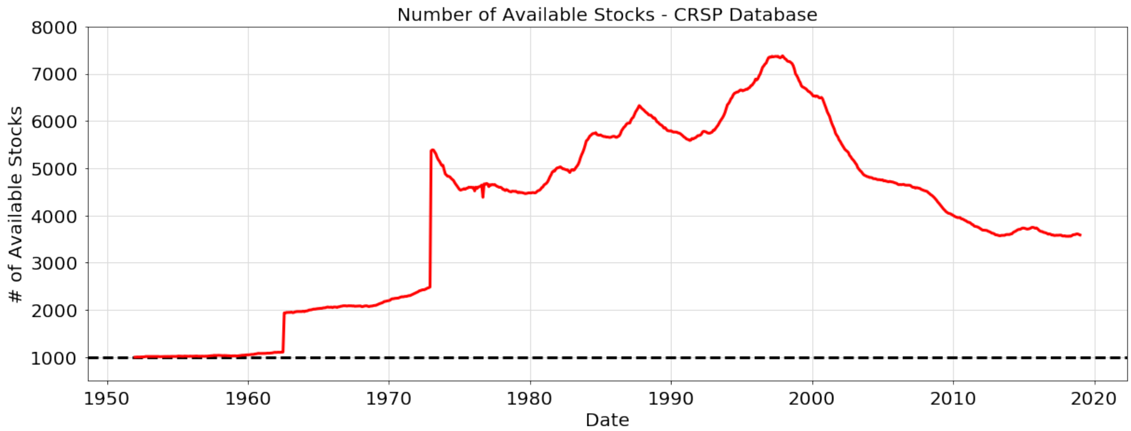

Figure 76: Number of stocks available in the CRSP database
The figure above represents the number of stocks available in the CRSP database each month which meet the selection criteria for this analysis. ETFs and ADRs are excluded from the full database as well as issues that CRSP does not classify as common equity shares. The dashed line represents the minimum number of 1,000 stocks needed for this analysis.

7.2.3 Exclusions

For the purpose of this analysis, all listings which are not classified as common stocks are excluded. This excludes listings such as exchange traded funds (ETFs), American Depository Receipts (ADRs), and other classifications which are not the common shares of corporations.

7.2.4 Delisting Returns

Monthly returns for all common shares are collected from the database. Some listings disappear from the database due to merger/acquisition, bankruptcy, delisting requirement, or other event. Since only full month returns are provided in the main database, when these events happen mid-month, the returns of that month are not reported in the main data set. For example, the bankruptcy of Lehman Brothers on September 15, 2008 would not have been fully recorded in the main data set.

Delisting month returns are provided by CRSP in a separate table. The returns in this analysis are enhanced with partial month delisting returns. While the overall effect is small, including these additional data points enriches the data set.

Additionally, I considered the work of (Shumway 1997) who documents a delisting bias in the stock return data of CRSP. While I have included the delisting returns provided, the author points out that there are many delistings for which returns are missing. By researching news feeds, Shumway creates a probable rate of return for the average missing delisting return. In addition, he also shows his analysis by assuming that all missing returns are replaced by a -100%. The effect on the overall data set is significant in that it creates a positive bias in the returns of the full data set. However, when the data set is sorted into size deciles, the impact is greatly diminished for the largest size deciles. The author concludes that "Including delisting returns of -1 reduces the performance of small stocks substantially, but it does not affect large stocks' returns". My analysis focuses only on the largest capitalization stocks available each year. For this reason, the assumptions about missing delisting returns is not included in my analysis.

7.3 Panel of Capital Allocation Methods (Weighting Schemes)

It has been well documented in previous studies such as (Chow et al. 2011) and (Clare, Motson, and Thomas 2013) that following any one of a range of alternative capital allocation methods would have led to superior returns and/or return efficiency. In these previous studies, the authors analyzed a wide range of methods and parameters. My study begins by selecting a panel of established and practical allocation methods and focuses on documenting the drivers of Sharpe ratios. To maintain focus on the interaction effect, a smaller set of allocation methods have been chosen which are more prevalent commercially and have wider practical appeal.

For purposes of clarity, the term used in this paper is 'capital allocation method' which has the same implied meaning as 'weighting scheme' which is a more common term in some parts of the world. In either form, this step in the portfolio formation process is concerned with spreading the available money across the securities which have been selected for the portfolio.

7.3.1 Market Capitalization

Market capitalization weighting is the most commonly available form of allocation method for investing passively and this is seen in the major market indices such as the S&P 500. While this allocation method is shown historically to have returns which are inferior to other allocation methods, it is highly efficient in a practical manner. As market capitalization of each security changes due to price movements each day, the weights change in unison requiring no trading turnover until the index is reformed.

The CRSP database provides share price and the number of shares outstanding each month. These two elements are multiplied together to create the market capitalization for each stock each period.

For consistency across the panel of allocation methods, I form a market capitalization weighted portfolio from the 1,000 largest securities available each year rather than use a market provided index which would have slightly different securities than would be used in my other allocation methods. This is an approach that is observed in the literature. The weighting for each stock (w_i) is found by dividing the individual stock's market capitalization by the sum of all stock's market capitalization on the date of index formation.

$$w_i = \frac{MktCap_i}{\sum_{i=1}^{N} MktCap_i} \tag{14}$$

The weights for the index are calculated using data from the last day of each year and the index is reformed and held without rebalance for the following year until the process is repeated. While this is not necessarily the exact selection and rebalance method that commercially available indices such as the Russell 1,000 implement, it is a method that can be applied consistently across a panel of allocation methods for this study and is consistent with methods employed in existing literature.

7.3.2 Equal Weight

In a well documented study, (DeMiguel, Garlappi, and Uppal 2009) show the benefits and simplicity of an equal weight allocation methodology. By following an equally weighted method, there are no parameters to estimate, and the weight calculation is trivial as it is simply the inverse of the number of stocks selected for inclusion in the index.

$$w_i = \frac{1}{N} \tag{15}$$

In addition to describing the simplicity and robustness of the equal weighted method, the authors show the period of historical data that would be required in order to accurately estimate parameters for other methods in order to outperform the simple equal weighted approach. A sample mean variance optimization based portfolio with 50 assets would require 6,000 monthly observations in order to accurately estimate the covariance matrix and outperform the simple equal weighted method.

One note of caution is provided by (Chow et al. 2011). The authors state that "A notable feature of equal weighting is that the resulting portfolio risk-return characteristics are highly sensitive to the number of included stocks. Although the S&P 500 Index and the Russell 1,000 Index have nearly identical risk-return characteristics over time, the equal-weighted S&P 500 portfolios and the equal-weighted Russell 1,000 portfolios have dramatically different risk-return characteristics."

7.3.3 Inverse Volatility

The inverse volatility allocation method attempts to distribute risk more equally across the set of available stocks while estimating a single parameter, volatility. As a result, one measure of diversification (effective number of stocks) is increased dramatically over the market capitalization method.

The method also attempts to capture a documented low-volatility anomaly whereby stocks with lower volatility have been shown to have higher risk adjusted returns. Building on their previous work from 1972, (Baker and Haugen 2012) show in a simple and transparent way that this anomaly has presented itself across long periods of time (US 1926 - 1990) and across the world (33 global stock markets during 1990 - 2011).

To create the weights for each stock in my study, the standard deviation of each stock (σ_i) is estimated using the previous 5 years of monthly returns. By using the inverse of this measure ($\frac{1}{\sigma_i}$), this allocation method shifts weight from stocks with higher volatility to stocks with lower volatility. The weight for each stock then is given by:

$$w_i = \frac{\frac{1}{\sigma_i}}{\sum_{i=1}^{N} \left(\frac{1}{\sigma_i}\right)} \tag{16}$$

The inverse volatility method combines increased diversification with harvesting the low-risk anomaly. While the method uses a single estimated parameter (standard deviation), the assumption under which this method would produce the optimal risk adjusted portfolio is that correlations are uniform across all securities.

7.3.4 Equal Risk Contribution

When the uniform correlations assumptions required for optimality of the inverse volatility method are not realized, then groups of highly correlated stocks receive larger allocations of risk. By including correlations between stocks into the allocation process, the equal risk contribution method accounts for the non-uniform correlation structure.

Estimating a second parameter introduces additional risks to the calculation. Calculating the weights which provide an equal contribution to risk from each stock does not have a closed form solution and must be solved using optimization. Optimization methods tend to amplify estimation errors, and this method now encounters estimation problems, especially as the covariance matrix is highly dimensional with up to 1,000 stocks.

To minimize the effect of these estimation errors, the analysis employs the shrinkage method described in (Ledoit and Wolf 2004) which condenses the outlier observations toward the central values in a systematic way and reduces the amplification of errors produced by the optimizer. The calculations are performed in Python using sklearn.covariance.LedoitWolf. Use of this shrinkage method is extended for all allocations in my research that require a covariance matrix.

All of the allocation methods examined in this study require long only positions ($w_i \geq 0$) with a constraint to be fully invested $\left(\sum_{i=1}^{N} w_i = 1 \right)$. To solve for the weights which produce an equal contribution to risk of each stock we follow the method described by (Maillard, Roncalli, and Teiletche 2008) and first define the minimization function.

Let w be the vector of weights and let Σ be the covariance matrix. Then the volatility of the portfolio is defined as:

$$\sigma(w) = \sqrt{w'\Sigma w} \tag{17}$$

The risk contribution of each asset within the portfolio is given as:

$$\sigma_i(w) = w_i \times \partial_{w_i} \sigma(w) \tag{18}$$

From this is constructed the vector of marginal contributions ($\partial_{w_i} \sigma(w)$) as:

$$c(w) = \frac{\Sigma w}{\sqrt{w'\Sigma w}} \tag{19}$$

[13] https://scikit-learn.org/stable/modules/generated/sklearn.covariance.LedoitWolf.html

The solution for the weights is given by the optimization:

$$\underset{w}{arg\,min} \sum_{i=1}^{N} \left[\frac{\sqrt{w'\Sigma w}}{N} - w_i \cdot c(w)_i \right]^2 \tag{20}$$

The minimization routine for this analysis was constructed with the Python 'SciPy' package using the sequential least squares programming (SLSQP) method.[14]

7.3.5 Minimum Variance

In a traditional mean-variance optimization approach, three parameter estimates are required: returns, volatilities, and correlations. Returns are the most difficult to forecast and errors in the estimated returns create large errors in an optimization output even when using a covariance shrinkage model. If it can be assumed that all stock returns are the same, then the optimal solution can be found by estimating just the volatility and correlation (as in the equal risk method) and solving for the weights that minimize the risk of the portfolio.

Depending on the parameter estimates, this allocation method may create a very concentrated set of weights. To prevent this possibility of concentration, my analysis follows the method of previous academic studies in limiting the maximum position size to 1%. For comparison purposes relative to the maximum diversification approach, I also include maximum position sizes at 5%. As above, this analysis requires long only positions ($w_i \geq 0$) and constrained to be fully invested ($\sum_{i=1}^{N} w_i = 1$).

To solve for the weights which produce a minimum variance portfolio, the minimization function is defined so that w is the vector of weights and Σ is the covariance matrix.

[14] `https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html`

The optimal weights are found by:

$$\underset{w}{min}(w'\Sigma w) \, \text{subject to} \begin{cases} \sum_{i=1}^{N} w_i = 1 \\ 0 \leq w_i \leq max\% \end{cases} \tag{21}$$

7.3.6 Maximum Diversification

Assuming that all stocks have the same return is unlikely to be true. (Choueifaty and Coignard 2008) propose an optimization method to maximize the diversification of a portfolio by instead assuming that returns are proportional to the risks of each stock. With R_f being the risk free rate, and the expected return given by $E(R_i)$, this relationship is given by:

$$E(R_i) - R_f = \gamma \sigma_i \tag{22}$$

where $\gamma > 0$ is the proportion of return per unit of risk.

The second major contribution of (Choueifaty and Coignard 2008) is the quantifiable measure of diversification known as the Diversification Ratio (DR).

$$DR = \frac{\sum_{i=1}^{N} w_i \sigma_i}{\sigma_p} \tag{23}$$

DR is the ratio of the weighted sum of the individual stock volatilities ($\sum_{i=1}^{N} w_i \sigma_i$) to the fully correlated and weighted sum of volatilities given by the portfolio volatility (σ_p).

Solving for the vector of weights that maximizes the diversification ratio is given by:

$$\max_{w}\left(\frac{w'\sigma}{\sqrt{w'\Sigma w}}\right) \text{subject to } \begin{cases} \sum_{i=1}^{N} w_i = 1 \\ 0 \leq w_i \leq max\% \end{cases} \tag{24}$$

For comparison purposes relative to the minimum variance approach, I use maximum position sizes of 1% and 5%. As can be seen in Table 17, even at these lower position size limits, the approach can lead to a small selection of securities for the portfolio.

The original work of (Choueifaty and Coignard 2008) was based on stocks in the S&P500 index from 1990 to 2008 and used 250 days of data to construct the parameter estimates required for the covariance matrix. Position limits were set at 10% but were additionally limited to month end weights which complied with UCITS III[15] rules. This restriction included a maximum limit of 40% on the sum of individual stock weights which exceed 5%. Even with these additional restrictions, setting position limits at this level often leads to very concentrated portfolios.

My simplified yet comparable approach (which does not include the UCITS restrictions) is calculated across the time frame of 1968 to 2017, from a universe of 1,000 stocks. This method on average carried an effective number of stocks of just 24. Using this method but with 10% position limits, I observe that although the methodology maximizes the ex-ante DR, the ex-post DR is only slightly above that of the market cap weighted method and on average below all of the other methods examined in this study. I do not include the 10% max position size version in my analysis but do note the effects of a high concentration risk as will be shown in my research when combining the steps of capital allocation with perfect stock selection.

[15] `https://www.esma.europa.eu/databases-library/interactive-single-rulebook/ucits`

7.4 Statistical Measures

7.4.0.1 Active Share: Active Share (AS) is the percentage of fund holdings that are different from the benchmark holdings. A fund that has no holdings in common with the benchmark will have an Active Share of 100%, and a fund that has exactly the same holdings with exactly the same weights as the benchmark considered will have an Active Share of 0%. With N being the number of stocks under consideration and $w_{fund,i}$ and $w_{index,i}$ being the weights of stock i, the formula for Active Share is given by:

$$AS = \frac{1}{2}\sum_{i=1}^{N}|w_{fund,i} - w_{index,i}| \tag{25}$$

The current description from the authors states that Active Share can most easily be calculated as 100% minus the sum of the overlapping portfolio weights. [16]

7.4.0.2 Annualized Volatility: This is the measure of standard deviation of returns. The figure is calculated using monthly returns and expanded to a yearly (annualized) measure by multiplying the calculation by the square root of the periodicity, in this case $\sqrt{12}$. r are the monthly return observations and T is the number of observed periods in a year. The formula is given by:

$$AnnVol = \sigma_r\sqrt{T} \tag{26}$$

7.4.0.3 Arithmetic Returns: Arithmetic returns are the average period returns without regard to sequence or compounding. R_a is the arithmetic return, N is the number of periods in the calculation, and r_i is the individual period return. The formula is

[16] https://activeshare.info/about-active-share

given by:

$$R_a = \frac{1}{N}\sum_{i=1}^{N} r_i = \frac{r_1 + r_2 + r_3 ... r_N}{N} \tag{27}$$

7.4.0.4 ArithRet / DaR: This is the simple ratio of average annual return relative to the 5% worst drawdown event observed. The metric captures a version of reward vs. risk.

7.4.0.5 Average Turnover: The turnover ratio is an indication of the trading activity within a portfolio. It is calculated by taking the sum of either the value of new purchases or securities sold (whichever is smaller) as a fraction of the average portfolio value. The turnover ratios in this analysis are an average of all of the individual year's turnover ratios.

7.4.0.6 Conditional Drawdown at Risk (CDaR 5%): CDaR is the average of the worst drawdowns observed that are worse than DaR. For example, CDaR 5% would give the mean drawdown that was observed in excess of the DaR 5% level.

7.4.0.7 Diversification Ratio(DR): The DR of a portfolio is defined as the ratio of the weighted average volatility of individual securities in that portfolio, divided by the volatility of the overall portfolio. The formula is given by:

$$DR = \frac{\sum_{i=1}^{N} w_i \sigma_i}{\sigma_p} \tag{28}$$

with w_i being the weight of a security, σ_i being the standard deviation of the returns of a security, and σ_p being the standard deviation of the returns of the portfolio.

7.4.0.8 Downside Volatility:

The downside volatility is a modification of the standard deviation concept where only observations below a minimum return are considered in the calculation. For the purpose of this analysis, the minimum return was set at zero. As with annualized volatility, the calculation is expanded at the square root of the observed periods in a year. The formula is given by:

$$DownVol = \frac{\sum_{i=1}^{N}(\min(r_i - 0, 0))}{N}\sqrt{T} \tag{29}$$

7.4.0.9 Drawdown at Risk (DaR 5%):

From a distribution of drawdown events, the DaR is the level of drawdown which contains the portion of events specified. For example, DaR 5% would be the level of drawdown at which 5% of the drawdown events observed are worse, and 95% of the observed drawdowns are better.

7.4.0.10 Effective Number of Stocks:

The Herfindahl-Hirschman index (HHI) measures the concentration ratio of a number of firms in an industry. The inverse of this index ($\frac{1}{HHI}$) gives a measure of diversification. The variable w_i is the percentage weight of each stock in the portfolio. The formula is given by:

$$\frac{1}{\sum_{i=1}^{N} w_i^2} \tag{30}$$

7.4.0.11 Geometric Returns:

Geometric returns are the geometric mean of a sequence of returns. Order of returns matters and the standard deviation of period returns has a negative effect when compared to arithmetic returns. This type of calculation is appropriate where returns exhibit serial correlation and is the correct method to calculate growth in capital from a particular sequence of returns.

R_g is the geometric return, N is the number of periods in the calculation, and r_i is a single period return. The formula for geometric return is given by:

$$R_g = \left(\prod_{i=1}^{N}(1 + r_i)\right)^{\frac{1}{N}} = \sqrt[N]{(1 + r_1)(1 + r_2)(1 + r_3)...(1 + r_N)} \tag{31}$$

7.4.0.12 Information Ratio (IR): The information ratio is a measure of returns in excess of a benchmark relative to the tracking error between the two returns. R_b is the benchmark return and TE (given below) is the tracking error between the two sets of returns. The IR of the benchmark portfolio is zero. The formula is given by:

$$IR = \frac{R_p - R_b}{TE} \tag{32}$$

7.4.0.13 Kurtosis: Kurtosis is a measure of the relative thickness in the tails of a distribution. The kurtosis measure is the fourth standardized moment of a distribution.

$$Kurtosis = \mathbb{E}\left[\left(\frac{X - \mu}{\sigma}\right)^4\right] = \frac{\mu^4}{\sigma^4} \tag{33}$$

Within the Python code used for this analysis, the kurtosis statistic is calculated from the Scipy library as scipy.stats.kurtosis. [17]

7.4.0.14 Maximum Drawdown(MDD): This statistic is the maximum loss from a peak to a trough of a portfolio, before a new peak is attained. The statistic is expressed as a percentage loss from the peak in value. P is the peak value attained, L is the lowest point subsequent to the peak before the peak value is recovered. The formula

[17] `https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.kurtosis.html`

is given by:

$$MDD = \frac{P - L}{P} \tag{34}$$

7.4.0.15 Modified Sharpe Ratio: This ratio was first suggested by (Favre and Galeano 2002) and (Gregoriou and Gueyie 2009) extending the original Sharpe ratio to account for the extreme observations in a return series and more fully capture the risk of the investment beyond the assumptions of a normal distribution embodied in the standard deviation. r_p is the return of the portfolio, r_f is the risk free rate (for this analysis: 3 month US Treasury Bill rate), $mVaR$ is the modified Value at Risk at a given confidence level (α), z_α is the z-score for the given level of confidence (-1.65 with 95% confidence), z_m is the z-score of a normal distribution for the given confidence level with a Cornish-Fisher expansion to account for skewness and excess kurtosis, s_p is skewness and k_p is kurtosis.

$$mSharpe(\alpha) = \frac{\overline{r_p - r_f}}{mVaR_p(\alpha)} \tag{35}$$

where

$$mVaR(\alpha) = \overline{r_p} - z_m(\sigma_p) \tag{36}$$

where

$$z_m = \left(z_\alpha + \frac{1}{6}(z_\alpha^2 - 1)s_p + \frac{1}{24}(z_\alpha^3 - 3z_\alpha)k_p - \frac{1}{36}(2z_\alpha^3 - 5z_\alpha)s_p^2 \right) \tag{37}$$

7.4.0.16 Sharpe Ratio: The Sharpe ratio is a measure of return efficiency per unit of risk. Returns above the risk free rate (excess return) are measured relative to the standard deviation of all returns. R_p is the return of the portfolio, R_f is the risk free rate (for this analysis: 3 month US Treasury Bill rate), σ_p is the annualized standard deviation of portfolio returns. The formula is given by:

$$Sharpe = \frac{R_p - R_f}{\sigma_p} \tag{38}$$

7.4.0.17 Skewness: Skewness is a measure of the asymmetry of the probability distribution of a variable around its mean. The skewness value can be positive or negative, or undefined. The skewness coefficient is the third standardized moment. X is a given observation, μ is the mean of the sample observations, σ is the standard deviation of the sample observations, and E is the expectation.

$$Skewness = \mathbb{E}\left[\left(\frac{X-\mu}{\sigma}\right)^3\right] = \frac{\mu^3}{\sigma^3} \tag{39}$$

7.4.0.18 Sortino Ratio: The Sortino ratio is a return vs. risk efficiency ratio and a variant of the Sharpe ratio. The return in the numerator are all returns above the minimum acceptable return (which is zero for this analysis) and the denominator is the downside volatility. The formula is given by:

$$Sortino = \frac{R_p - 0}{\sigma_{p(Downside)}} \tag{40}$$

Within the Python code used for this analysis, the skewness statistic is calculated from the Scipy library as scipy.stats.skew. [18]

[18] `https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.skew.html`

7.4.0.19 Tracking Error(TE):

Tracking error is the standard deviation of the difference in returns between a portfolio and its benchmark. The formula is given by:

$$TE = \sqrt{\frac{\sum_{i=1}^{N} (R_p - R_b)^2}{N - 1}} \tag{41}$$

7.5 Comparison of Capital Allocation Methods

The analysis begins by detailing the performance and statistical attributes of the panel of allocation methods using the full 1,000 largest stocks available in the data set each year.

7.5.1 Statistical Characteristics

Using the data and capital allocation methods described above, portfolios are formed annually at the end of each calendar year. From the return series and portfolio holdings for each method, a panel of statistics is calculated and shown in Table 17. The calculations for these statistical measures are fully described in the appendix of this research.

7.5.1.1 Returns:

The metric of return is a measurement of the payment that investors earn over a period of time, presumably for the assumption of some risk. Any estimate of this metric will be imperfect due to issues such as autocorrelation, mean reversion at long horizons, et cetera. In addition, observations of stock returns over various horizons depict a distribution which is not normal. Therefore the calculation of an unbiased estimate of return is problematic.

In an early paper, (Blume 1974) documented the biases that exist with both arithmetic

Table 17: Summary statistics for various capital allocation methods, 1968 - 2017

	MC	**EW**	**IV**	**ER**	**MV1%**	**MV5%**	**MD1%**	**MD5%**
Geom Ret %	10.2	11.7	12.2	12.0	11.6	11.2	11.4	10.9
Arith Ret %	11.4	13.3	13.5	13.3	12.4	11.9	12.4	12.0
Skewness	-0.39	-0.35	-0.33	-0.38	-0.40	-0.66	-0.60	-0.65
Excess Kurtosis	1.77	2.16	2.46	2.54	3.63	2.72	2.44	2.47
Ann Vol %	15.1	17.0	15.5	15.2	11.9	11.3	13.8	14.3
Down Vol %	9.6	10.7	9.6	9.4	7.3	7.0	8.8	9.3
Track Error %		5.4	4.9	4.8	7.2	8.0	6.3	7.6
Max DD %	-50.5	-52.2	-49.5	-48.9	-38.6	-36.8	-42.7	-45.5
DaR 5%	-33.7	-27.6	-26.2	-25.5	-19.9	-19.1	-21.5	-22.3
CDaR 5%	-45.1	-42.9	-44.0	-43.5	-35.1	-33.7	-34.8	-37.7
Sharpe Ratio	0.44	0.50	0.56	0.56	0.64	0.63	0.55	0.50
Modified Sharpe	0.17	0.19	0.21	0.20	0.22	0.21	0.20	0.18
Sortino Ratio	0.70	0.80	0.92	0.91	1.06	1.03	0.88	0.79
Info Ratio		1.88	2.06	2.1	1.4	1.3	1.6	1.3
Ret / DaR	0.32	0.45	0.48	0.49	0.59	0.59	0.55	0.51
Effect # Stocks	138	1,000	877	831	113	52	113	44
Divers Ratio	2.04	2.16	2.13	2.21	2.33	2.35	2.65	2.80
Avg Turnover %	8.8	13.0	12.5	13.9	36.8	43.2	40.3	47.0
Active Share		47.1	49.2	52.1	84.8	90.1	88.0	93.6

MC: 'Market Capitalization' holds all 1,000 stocks and weights the holdings annually based on the relative market capitalization of each stock to the total market capitalization of all stocks on that date. **EW:** 'Equal Weight' holds all 1,000 stocks and weights each stock equally. **IV:** 'Inverse Volatility' gives more weight to stocks with lower volatility and less weight to stocks with higher volatility as measured monthly over the previous five year. **ER:** 'Equal Risk' holds all 1,000 stocks, weights solved by optimization method whereby the marginal contribution of each stock to total portfolio risk is equalized. **MV1%/5%:** 'Minimum Variance' holds a subset of the available stocks where weighting is based on an optimization which finds the portfolio with the minimum total variance, subject to maximum position sizes of 1% or 5%. **MD1%/5%:** 'Maximum Diversification' holds a subset of stocks each year with weights based on an optimization which maximizes the ex-ante diversification ratio (DR) of the portfolio using volatility and correlation data monthly from the previous five years. Position sizes are limited to 1%, 5%. **Statistical calculations contained in this table are explained in the appendix of this chapter.**

and geometric returns. In general, arithmetic returns are biased to the upside while geometric returns are biased to the downside. The author developed four unbiased estimators of return and showed their efficiency. In a subsequent paper, (Jacquier, Kane, and Marcus 2003) showed that the unbiased estimate is a weighted average between arithmetic and geometric returns. For short periods, arithmetic is heavily weighted and as the horizon return begins to approach the estimation period data, geometric is most heavily weighted.

Advanced methods for correcting for these biases and designing an efficient estimator was proposed by (Jacquier 2005).

An arithmetic return estimate is constructed as the simple mean of the observed single period returns. Geometric returns are calculated as the single period return that would allow an investor to earn the total compounded return observed over the sample period. For the analysis in this dissertation, both metrics are calculated and any biases that would appear from using these estimates for forecasting future returns are inherent symmetrically across each of the allocation methods. With a more thorough basis for understanding return estimates, the focus now turns to a comparative look across the return statistics of each allocation method.

In general we see that both geometric and arithmetic returns for the market capitalization approach are the lowest while those of inverse volatility and equal risk are the highest. All alternative allocation methods outperform the market capitalization approach for both geometric and arithmetic return measures for the 50 years of this study. For the purpose of this research I focus on the arithmetic returns rather than geometric returns. Each has particular merits and restrictions.

7.5.1.2 Return Distribution Normality: Distributions with negative skew have more observations in the left tails than would be implied by the normal distribution. This is a bad characteristic for investors and should require a higher return premium as compensation for this risk. Skewness for the market capitalization and the three allocation methods (EW, IV, ER) that contain all 1,000 stocks are relatively well behaved.

But for the minimum variance and the maximum diversification approaches, negative skewness is more pronounced with the worst observations being at the 5% maximum position size for each of these optimization methods.

Kurtosis measures the weight in both tails together. A normal distribution has a kurtosis of 3. This analysis shows excess kurtosis or the amount in excess of a normal distribution. The market capitalization approach has the lowest observed excess kurtosis at 1.77. By contrast, the minimum variance methods shows the highest levels of excess kurtosis.

The combination of excess kurtosis and negative skew are especially bad and would naturally be expected to warrant a higher return premium as compensation. The worst combinations of excess kurtosis and skew are observed in the minimum variance methods. While the additional premium required for this combination is not seen simply in the returns of these approaches, they do show the highest ratios of reward vs. risk such as Sharpe, Sortino, and Return/DaR which is consistent with this expectations. The Modified Sharpe ratio (explained below) accounts for these elements of non-normality in the distributions of returns.

7.5.1.3 Volatility: The volatility for the market cap weighted portfolio sits in the middle of all of the other approaches. Equal weight has the highest volatility which may be due to the size bias that would naturally exist in a method which spreads the allocation weights equally relative to an approach which is highly concentrated in the largest capitalization names. As would be expected, the minimum variance approaches show the lowest realized volatility but this is clearly a function of the metric that is optimized during the construction of the portfolio.

7.5.1.4 Drawdown: The are three metrics to examine in this area. All values are calculated on a monthly basis and indicate the loss from a previous peak in value.

Maximum drawdown is the single largest loss event experienced over the entire history. While most of the other allocation methods show improvements in this area, the equal weight method has larger losses during the worst single event. The greater size of drawdowns seen in the market cap method are likely functions of the effective concentration of these portfolios into a very few stocks. The large drawdowns from the equal weight method may be explained by the increased exposure to smaller cap securities.

Drawdown at Risk (DaR 5%) examines the distribution of all drawdown events and denotes the size of the 5% worst event. It is important to note that each complete event (peak to trough to peak) is treated as a single observation for the distribution. Once we move beyond examining just the single largest drawdown event, the DaR 5% tells a different story. Every other allocation method has better metrics with improvements in the range of 6.1% to 14.6% better than MC. During a relatively normal sized drawdown event, any method other than the market capitalization allocation method performed better across this time frame.

The final metric is the **Conditional Drawdown at Risk** (CDaR 5%) which examines the average size of the drawdown which is in excess of the DaR 5%. CDaR 5% calculates the mean loss that occurs when the loss is within the left tail. Maximum drawdown is the single worst observation in the distribution and it will have a direct impact on the CDaR calculation. It is interesting to note that despite the equal weight method having larger absolute maximum drawdowns, the market capitalization method still exhibits the highest CDaR 5%.

7.5.1.5 Reward vs. Risk: For both the Sharpe ratio and Sortino ratio, the market capitalization method shows the lowest mean statistics. As noted above, the minimum variance (1% max position & 5% max position) methods appear to be clear winners for nominal risk adjusted returns, but this will be challenged with robustness testing. These higher risk adjusted returns may be the required compensation for the combination of negative skew and excess kurtosis that these strategies exhibit as the Modified Sharpe Ratio (below) for each of these methods are in line with the other allocation methods. All other allocation methods appear to outperform the market capitalization method, but this will be tested.

The **Sharpe Ratio** was originally proposed by (Sharpe 1966) as the reward to variability ratio. In it's original form, the equation given was:

$$R/V = \frac{A_i - p}{V_i} \tag{42}$$

with A_i being the return of the investment, p being the risk free rate, and V_i being the standard deviation of the returns on investment i. A_i was described by Sharpe as the 'average annual rate of return'. Unfortunately, there is no single statistical measure known as 'average' and so I default to the most simple approach which is the arithmetic average of returns. Practitioners sometimes use geometric averages.

The Sharpe ratio is a common measure of risk adjusted returns, however as the denominator is simply the standard deviation of the return series, it suffers from being indiscriminate relative to distributions with high levels of skew and excess kurtosis which are generally disliked by investors.

The **Modified Sharpe** was originally suggested by (Favre and Galeano 2002) and then by (Gregoriou and Gueyie 2009). The Modified Sharpe ratio replaces the denominator of the original Sharpe ratio (standard deviation of returns) with a modified Value at Risk (mVaR) calculation which accounts for the extreme observations more fully by incorporating skew and excess kurtosis. With a standard parametric VaR measure, a z-score is used as a scalar on the standard deviation of the observations. The mVaR calculation uses a Cornish Fisher expansion to adjust the z-score for the non-normality of the distribution before scaling the standard deviation observation.

The mVaR ratios are not directly comparable to standard Sharpe ratios. They are however more comparable to each other as each has been properly adjusted for the first four moments of their distributions.

A case in point can be made by looking at the ratios for the IV and ER allocation methods. Standard Sharpe ratios are both 0.52 while the modified Sharpe ratios are different. As ER has slightly more negative skew and a higher excess kurtosis, the modified Sharpe is slightly lower with 0.21 (IV) and 0.20 (ER). In cases where distributions have extreme third and fourth moments, the modified Sharpe is key to adjusting these ratios and leveling the comparisons for risk adjusted returns.

The **Information Ratio** is highest with the inverse volatility and equal risk contribution methods which is a function of the low tracking error that these method exhibit. As tracking error is the denominator of the information ratio, the methods with the highest tracking errors (minimum variance and maximum diversification) have the lowest information ratios. These two methods have the largest differences in construction as will be shown with Active Share and while there are risk adjusted performance advantages to straying from the benchmark, the excess tracking risk and information ratios are negatively impacted by these approaches.

7.5.1.6 Diversification: There are several ways to consider diversification. The most simplistic approach is simply the number of stocks in the portfolio. The market capitalization, equal weight, inverse volatility, and equal risk methods the portfolios hold all 1,000 stocks that are available on each date. The minimum variance and maximum diversification portfolios hold a smaller subset of stocks due to the optimization routine preferring some stocks to carry zero weight.

Another measure of diversification to consider is the **Effective Number of Stocks** which is a function of the allocation weights. This metric is calculated as the inverse of the Herfindahl-Hirschmann Index (HHI)[19] which is a measure of concentration.

$$HHI = \sum_{i=1}^{N} w_i^2 \tag{43}$$

The inverse of the HHI is the Effective Number of Stocks. Large weights to a few stocks will make this ratio smaller while low weights to many stocks will make this ratio larger.

Despite holding 1,000 stocks, the market capitalization method is only effectively influenced by 138 stocks. This is in stark contrast to the equal weight method which holds the same exact stocks but in equal proportion so that the effective number of stocks held is exactly 1,000. Similarly, the inverse volatility and the equal risk portfolios have very high effective number of stocks and would appear to be more diversified than the market capitalization method, even though they all own the same set of stocks each year. With a smaller set of stocks held and relatively large concentration weights, the minimum variance and maximum diversification methods display a metric for effective number of stocks that is less than the market capitalization method.

A third method to consider is the **Diversification Ratio** which is a function of weights, volatilities, and the correlations between securities. While the maximum diversification method seeks to maximize this ratio ex-ante, the ratio presented in Table 17 on

[19] `https://corporatefinanceinstitute.com/resources/knowledge/finance/herfindahl-hirschman-index-hhi/`

page 237 is the ex-post result. Again we see the market capitalization method being the least diversified and all other allocation methods are superior in this ratio. As noted previously but not included in our analysis, while the maximum diversification (10% max position) is the least constrained version which attempts to maximize this ratio, it barely shows any difference relative to the most concentrated method which is the market capitalization. There appears to be a wide spread between the optimization of this ratio ex-ante and the ex-post resulting DR when position sizes are allowed to range quite high.

7.5.1.7 Turnover: The market capitalization method has the lowest turnover at 8.8% per year on average. This is unsurprising as the weights of stocks in the portfolio float with changes in value and are only influenced by names being added or deleted from the universe once per year. The second lowest set of turnover is seen with the three methods that hold the same set of 1,000 stocks but have their weights changed independently from the value of the stocks. These are the equal weight, inverse volatility, and equal risk contribution methods. Finally, the minimum variance and maximum diversification methods change the subset of stocks held and solve for a new set of weights each period which drives their average annual turnover much higher than the other methods.

A final metric shown was put forth by (Cremers and Petajisto 2009). **Active Share (AS)** uses the current holdings to measure the 'active bets' that a manager is making relative to the benchmark. The authors explain the measure as 100 minus the overlapping position weights. AS can increase as stocks are either included or excluded and as the weights of stocks vary relative to their weights in the benchmark.

"Active management" tends to refer to managers who engage in security selection. However, according to ranges of AS set by the authors, all of the alternative allocation methods in this study show AS scores in the range of active management even though none of them are engaging in intentional stock selection. Equal weight, inverse volatility, and equal risk contribution hold all of the same stocks as the benchmark but only vary their weights. Minimum variance and maximum diversification hold a subset of stocks and so the higher AS measure are a function of both of these variables. Minimum variance and maximum diversification methods show AS scores in the range of 85-95.

7.5.2 Factor Exposures

7.5.2.1 Fama-French 3 Factor Model (FF3): This asset pricing model, developed by (Fama and French 1992), extends the original single factor capital asset pricing model (CAPM) by including exposure sensitivity to firm size and valuation ratios. These additional factors have been shown to provide greater explanation for individual equity returns and help to increase the R^2 of the regression.

Exposures (betas) for the monthly returns of each portfolio allocation method are shown in Table 18 relative to the factors of market, size, and value as provided by Kenneth French[20]. Calculations are based on a multiple regression analysis and calculated in Python using the Statsmodels[21] package which is documented by (Seabold and Perktold 2010).

It is important to note that although I have constructed a market capitalization weighted portfolio as the benchmark for our analysis, our methodology differs slightly from that used for the creation of the factor returns. Stock inclusion may be slightly different, and we rebalance annually while the factor return portfolios are rebalanced quarterly.

- **Intercept:** The intercept in the equation is the additional performance not ex-

[20] `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`
[21] `https://www.statsmodels.org/stable/index.html`

Table 18: Fama-French Three Factor Model, 1968 - 2017

	MC	EW	IV	ER	MV1	MV5	MD1	MD5
Intercept	0.00%	-0.02%	0.00%	0.00%	0.12%	0.12%	0.07%	0.04%
t-stat	-0.07	-0.47	0.67	0.85	1.94	1.78	1.12	0.54
Market	0.99	1.04	0.96	0.94	0.71	0.65	0.81	0.79
t-stat	216.87	102.22	90.40	91.57	48.81	41.54	55.88	43.67
SMB	-0.12	0.27	0.17	0.17	0.04	0.05	0.18	0.22
t-stat	-17.94	18.75	10.93	11.60	1.88	2.15	8.47	8.36
HML	0.06	0.27	0.33	0.30	0.27	0.25	0.21	0.20
t-stat	8.16	17.62	20.05	18.90	12.19	10.60	9.69	7.36
R^2	0.99	0.96	0.94	0.94	0.81	0.76	0.86	0.80

Monthly excess returns above the risk free rate for each allocation method are regressed against the following set of standard factors as collected from: `http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/index.html`. The website contains full information on the formation of the factor return sets. **Market:** excess return above the 1 month US Treasury Bill of the value weighted return of all CRSP firms incorporated in the US and listed on NYSE, AMEX, or NASDAQ exchanges with a share code of 10 or 11. **SMB:** Small Minus Big - the average return on the three small portfolios minus the return of the three big portfolios. **HML:** High Minus Low - the average return on the two value portfolios minus the average return on the two growth portfolios. Capital allocation methods: MC - market capitalization, EW - equal weight, IV - inverse volatility, ER - equal risk, MV1 - minimum variance 1% max position size, MV5 -minimum variance 5% max position size, MD1 - maximum diversification 1% max position size, MD5 - maximum diversification 5% max position size

plained by the combination of factor exposures solved in the multiple regression. Using a factor lens, the intercept is the unexplained performance alpha (α). Given the performance improvements that are seen in the gross return numbers for the alternative allocation methods, the difference in the intercepts relative to our version of the market cap portfolio are quite small as most of the performance is explained by combining market returns with a few additional factors.

The t-statistic is shown for each value of the intercept. Our null hypothesis is that the true alpha of any of these allocation methods is zero. Using a two tailed test at the 95% confidence level, this would make the critical value equal to 1.96 which would cause us to reject the alphas for all of the methods although the alpha for MV1 is very close to being significant.

Minimum variance portfolios are in several characteristic ways the most different than other allocation methods and have the largest intercepts although we have to reject these values at the 95% confidence level and assume that they are possibly zero.

- **Market:** The market exposure factor is the excess return of a value weighted portfolio of stocks above the risk free rate. For the three methods which hold all 1,000 stocks, the market exposure is quite high and very near that of the market capitalization method in our panel. As minimum variance and maximum diversification hold fewer stocks and have dramatically different weightings, the portion of returns attributable to the overall market are significantly lower. All of the t-statistics are significant for the market coefficients.

- **Size (SMB):** The size factor captures the return of small versus big stocks (market capitalization) within a portfolio. A positive value suggests more exposure to smaller capitalization stocks than in the market portfolio. As expected, all of the allocation methods examined in this study show more exposure to smaller capitalization stocks. This is by design as the market cap weighted portfolio concentrates its exposure in the largest stocks available. Of the methods that hold all 1,000 stocks each year, equal weight has the most exposure to smaller cap stocks. All of the t-stats are significant with the exception of MV1 which is close to the threshold.

- **Value (HML):** The value factor captures the excess return of cheap stocks relative to expensive stocks. This is measured through the book to market ratio. The full panel of allocation methodologies display additional exposure to the value factor in the range of 0.20 - 0.33. This is reasonably significant with all of the t-stats above the threshold limit of 1.96.

- **R Squared:** The allocation methods which hold the full 1,000 stocks (equal weight, inverse volatility, equal risk) are very well explained by the factors in the model with a minimum R^2 of 0.94. While the factor model explains less of the optimized methods (minimum variance, maximum diversification) returns, the R^2 is still relatively high. It is noticeable that the more concentrated strategies (max position sizes of 5%) have lower explanatory power by the factor model.

7.6 Computational Intensity

A quick note is warranted with regard to computational intensity of the methods examined in this paper. For the market cap, equal weight, and inverse volatility methods, solving for the weights for each security is trivial. For the equal risk, minimum variance, and maximum diversification methods, the solution requires an optimization routine which was implemented in Python with sequential least squares programming. The optimization routines are computationally intensive because the covariance matrix contains one million cells.

For my analysis I used an Amazon virtual machine with 48 cores and 192 GB of RAM. For the simple methods, I was able to create 50 years of returns within 1 minute. However for the advanced methods, with basic (non-optimized) coding, the optimization routines would regularly take up to 8 minutes per year, per method with some solutions extending as long as 8 hours for a 50 year series of returns[22]. I assume that with code developed more efficiently the speed of calculation would increase, but I make this note for posterity about the current state of computing power and my experience crafting the solutions for this research. Despite my use of advanced cloud based computing re-

[22]Optimization routines were constructed using Python SciPy sequential least squares programming where tolerances were set to $1e - 08$.

sources, recent announcements from Google about their ability to demonstrate quantum supremacy in computing will likely make the calculations in this research trivial in the years ahead.

References

Amenc, N., F. Goltz, and A. Lodh (2016). "Smart Beta Is Not Monkey Business". *The Journal of Index Investing* 6.4, pp. 12–29.

Amenc, N., F. Goltz, and A. Lodh (2012). "Choose Your Betas: Benchmarking Alternative Equity Index Strategies". *The Journal of Portfolio Management* 39.1, pp. 88–111.

Ang, A. (2014). *Asset Management.* 1st ed. Oxford University Press.

Ang, A. et al. (2006). "The Cross-Section of Volatility and Expected Returns". *The Journal of Finance* 61.1, pp. 259–299.

Arnott, R., J. Hsu, and P. Moore (2007). "Fundamental Indexation". *Financial Analysts Journal* 61.2, pp. 83–99.

Arnott, R. et al. (2010). "Beyond cap weighted The empirical evidence for a diversified beta". *Journal of Indexes* February.

Arnott, R. et al. (2013). "The Surprising Alpha From Malkiel'sMonkey and Upside-Down Strategies". *The Journal of Portfolio Management* 39.4, pp. 91–105.

Baker, M. et al. (2010). "Can Mutual Fund Managers Pick Stocks? Evidence from Their Trades Prior to Earnings Announcements". *Journal of Financial and Quantitative Analysis* 45.05, pp. 1111–1131.

Baker, N. and R. Haugen (2012). "Low Risk Stocks Outperform within All Observable Markets of the World". *SSRN Electronic Journal.*

Baks, K., A. Metrick, and J. Wachter (2001). "Should Investors Avoid All Actively Managed Mutual Funds? A Study in Bayesian Performance Evaluation". *The Journal of Finance* 56.1, pp. 45–85.

Barras, L., O. Scaillet, and R. Wermers (2010). "False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas". *The Journal of Finance* 65.1, pp. 179–216.

Bera, A. and S. Park (2008). "Optimal Portfolio Diversification Using the Maximum Entropy Principle". *Econometric Reviews* 27.4-6, pp. 484–512.

Berk, J. and J. van Binsbergen (2015). "Measuring skill in the mutual fund industry". *Journal of Financial Economics* 118.1, pp. 1–20.

Blume, M. E. (1974). "Unbiased estimators of long-run expected rates of return". *Journal of the American Statistical Association.*

Bollen, N. and J. Busse (2001). "On the Timing Ability of Mutual Fund Managers". *The Journal of Finance* 56.3, pp. 1075–1094.

Brinson, G., R. Hood, and G. Beebower (1995). "Determinants of Portfolio Performance". *Financial Analysts Journal* 51.1, pp. 133–138.

Brown, S. and W. Goetzmann (1997). "Mutual fund styles". *Journal of Financial Economics* 43.3, pp. 373–399.

Carhart, M. (1997). "On Persistence in Mutual Fund Performance". *The Journal of Finance* 52.1, p. 57.

Carvalho, R. Leote de, X. Lu, and P. Moulin (2012). “Demystifying Equity Risk–Based Strategies: A Simple Alpha plus Beta Description”. *The Journal of Portfolio Management* 38.3, pp. 56–70.

Chaves, D. et al. (2011). “Risk Parity Portfolio vs. Other Asset Allocation Heuristic Portfolios”. *The Journal of Investing* 20.1, pp. 108–118.

Chen, Z. and P. Knez (1996). “Portfolio Performance Measurement: Theory and Applications”. *Review of Financial Studies* 9.2, pp. 511–555.

Choueifaty, Y. and Y. Coignard (2008). “Toward Maximum Diversification”. *The Journal of Portfolio Management* 35.1, pp. 40–51.

Choueifaty, Y., T. Froidure, and J. Reynier (2011). “Properties of the Most Diversified Portfolio”. *SSRN Electronic Journal.*

Chow, T.-m. et al. (2011). “A Survey of Alternative Equity Index Strategies”. *Financial Analysts Journal* 67.5, pp. 37–57.

Clare, A., N. Motson, and S. Thomas (2013). “An Evaluation of Alternative Equity Indices - Part 1: Heuristic and Optimised Weighting Schemes”. *SSRN Electronic Journal* March, pp. 1–41.

Clarke, R., H. de Silva, and S. Sapra (2011). “Toward More Information-Efficient Portfolios”. *Modern Portfolio Management.* Hoboken, NJ, USA: John Wiley & Sons, Inc., pp. 323–341.

Clarke, R., H. de Silva, and S. Thorley (2012). “Risk Parity, Maximum Diversification, and Minimum Variance: An Analytic Perspective”. *SSRN Electronic Journal.*

Cremers, M. (2017). "Active Share and the Three Pillars of Active Management: Skill, Conviction, and Opportunity". *Financial Analysts Journal* 73.2, pp. 61–79.

Cremers, M. and A. Pareek (2016). "Patient capital outperformance: The investment skill of high active share managers who trade infrequently". *Journal of Financial Economics* 122.2, pp. 288–306.

Cremers, M. and A. Petajisto (2009). "How Active Is Your Fund Manager? A New Measure That Predicts Performance". *Review of Financial Studies* 22.9, pp. 3329–3365.

Cremers, M. et al. (2018). *Challenging the Conventional Wisdom on Active Management: A Review of the Past 20 Years of Academic Literature on Actively Managed Mutual Funds*. Tech. rep.

Cuthbertson, K., D. Nitzsche, and N. O'Sullivan (2008). "UK mutual fund performance: Skill or luck?" *Journal of Empirical Finance* 15.4, pp. 613–634.

Daniel, K. et al. (1997). "Measuring Mutual Fund Performance with Characteristic-Based Benchmarks". *The Journal of Finance* 52.3, p. 1035.

DeMiguel, V., L. Garlappi, and R. Uppal (2009). "Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?" *Review of Financial Studies* 22.5, pp. 1915–1953.

DeMiguel, V. and F. Nogales (2009). "Portfolio Selection with Robust Estimation". *Operations Research* 57.3, pp. 560–577.

Eling, M. and F. Schuhmacher (2007). "Does the choice of performance measure influence the evaluation of hedge funds?" *Journal of Banking and Finance* 31.9.

Elton, E., M. Gruber, and C. Blake (2001). "A first look at the accuracy of the CRSP Mutual Fund Database and a comparison of the CRSP and Morningstar mutual fund databases". *Journal of Finance* 56.6, pp. 2415–2430.

Fabozzi, F., D. Huang, and G. Zhou (2010). "Robust portfolios: contributions from operations research and finance". *Annals of Operations Research* 176.1, pp. 191–220.

Fama, E. (1972). "Components of Investment Performance". *The Journal of Finance* 27.3, pp. 551–567.

Fama, E. and K. French (1992). "The Cross-Section of Expected Stock Returns". *The Journal of Finance* 47.2, p. 427.

– (2010). "Luck versus Skill in the Cross-Section of Mutual Fund Returns". *The Journal of Finance* 65.5, pp. 1915–1947.

Favre, L. and J.-A. Galeano (2002). "Mean-Modified Value-at-Risk Optimization with Hedge Funds". *The Journal of Alternative Investments* 5.2, pp. 21–25.

Fernholz, R., R. Garvy, and J. Hannon (1998). "Diversity-Weighted Indexing". *The Journal of Portfolio Management* 24.2, pp. 74–82.

Frazzini, A., J. Friedman, and L. Pomorski (2016). "Deactivating Active Share". *Financial Analysts Journal* 72.2, pp. 14–21.

Fulkerson, J. A. and T. B. Riley (2019). "Portfolio concentration and mutual fund performance". *Journal of Empirical Finance* 51.

Glode, V. (2011). "Why mutual funds 'underperform'". *Journal of Financial Economics* 99.3, pp. 546–559.

Goetzmann, W. et al. (2007). “Portfolio Performance Manipulation and Manipulation-proof Performance Measures”. *Review of Financial Studies* 20.5, pp. 1503–1546.

Goyal, A. and P. Santa-Clara (2003). “Idiosyncratic Risk Matters!” *The Journal of Finance* 58.3, pp. 975–1007.

Gregoriou, G. N. and J.-P. Gueyie (2009). “Risk-Adjusted Performance of Funds of Hedge Funds Using a Modified Sharpe Ratio”. *The Journal of Wealth Management*.

Grinblatt, M. and S. Titman (1989). “Portfolio Performance Evaluation: Old Issues and New Insights”. *Review of Financial Studies* 2.3, pp. 393–421.

– (1994). “A Study of Monthly Mutual Fund Returns and Performance Evaluation Techniques”. *The Journal of Financial and Quantitative Analysis* 29.3, p. 419.

Hsu, J. (2006). “Cap-Weighted Portfolios are Sub-Optimal Portfolios”. *Journal of Investment Management* 4.3, pp. 1–10.

Huang, J., C. Sialm, and H. Zhang (2011). “Risk Shifting and Mutual Fund Performance”. *Review of Financial Studies* 24.8, pp. 2575–2616.

Hwang, I., S. Xu, and F. In (2018). “Naive versus optimal diversification: Tail risk and performance”. *European Journal of Operational Research* 265.1, pp. 372–388.

Ibbotson, R. (2010). “The importance of asset allocation”. *Financial Analysts Journal*.

Ibbotson, R. and P. Kaplan (2000). “Does Asset Allocation Policy Explain 40, 90, or 100 Percent of Performance? Authors’ Response”. *Financial Analysts Journal* 56.3, pp. 16–19.

Jacquier, E. (2005). “Optimal Estimation of the Risk Premium for the Long Run and Asset Allocation: A Case of Compounded Estimation Risk”. *Journal of Financial Econometrics* 3.1, pp. 37–55.

Jacquier, E., A. Kane, and A. J. Marcus (2003). “Geometric or arithmetic mean: A reconsideration”. *Financial Analysts Journal*.

Jensen, M. (1968). “The Performance of Mutual Funds in the Period 1945-1964”. *Journal of Finance* 23.2, pp. 389–416.

Jiang, G., T. Yao, and T. Yu (2007). “Do mutual funds time the market? Evidence from portfolio holdings”. *Journal of Financial Economics* 86.3, pp. 724–758.

Jobson, J. D. and B. Korkie (1981). “Performance Hypothesis Testing with the Sharpe and Treynor Measures”. *The Journal of Finance* 36.4, pp. 889–908.

Kacperczyk, M., S. V. Nieuwerburgh, and L. Veldkamp (2014). “Time-Varying Fund Manager Skill”. *The Journal of Finance* 69.4, pp. 1455–1484.

Kacperczyk, M., C. Sialm, and L. Zheng (2008). “Unobserved Actions of Mutual Funds”. *Review of Financial Studies* 21.6, pp. 2379–2416.

Kolm, P., R. Tütüncü, and F. Fabozzi (2014). “60 Years of portfolio optimization: Practical challenges and current trends”. *European Journal of Operational Research* 234.2, pp. 356–371.

Kosowski, R. et al. (2006). “Can Mutual Fund “Stars” Really Pick Stocks? New Evidence from a Bootstrap Analysis”. *The Journal of Finance* 61.6, pp. 2551–2595.

Ledoit, O and M Wolf (2004). “Honey, I shrunk the sample covariance matrix - Problems in mean-variance optimization”. *Journal of Portfolio Management* 30.4, pp. 110–119.

Ledoit, O. and M. Wolf (2008). “Robust performance hypothesis testing with the Sharpe ratio”. *Journal of Empirical Finance* 15.5, pp. 850–859.

Lintner, J. (1965). “Security Prices, Risk, and Maximal Gains from Diversification”. *The Journal of Finance* 20.4, pp. 587–615.

Lo, A. W. (2002). “The Statistics of Sharpe Ratios”. *Financial Analysts Journal* 58.4.

Maillard, S., T. Roncalli, and J. Teiletche (2008). “On the Properties of Equally-Weighted Risk Contributions Portfolios”. *SSRN Electronic Journal*, pp. 1–23.

Malkiel, B. G. (1995). “Returns from Investing in Equity Mutual Funds 1971 to 1991”. *The Journal of Finance* 50.2, pp. 549–572.

Markowitz, H. (1952). “Portfolio Selection”. *The Journal of Finance* 7.1, pp. 77–91.

Martellini, L. (2008). “Toward the Design of Better Equity Benchmarks”. *The Journal of Portfolio Management* 34.4, pp. 34–41.

Memmel, C. (2003). “Performance Hypothesis Testing with the Sharpe Ratio”. *Finance Letters* 1.1.

Merton, R. (1980). *On Estimating the Expected Return on the Market: An Exploratory Investigation*. Tech. rep.

Michaud, R. (2014). “The Markowitz Optimization Enigma: is ’Optimized’ Optimal?” *SSRN Electronic Journal*.

Modigliani, F. and L. Modigliani (1997). “Risk-adjusted performance: How to measure it and why”. *Journal of Portfolio Management* 23.2.

Pastor, R. Stambaugh, and L. Taylor (2017). “Do Funds Make More When They Trade More?” *The Journal of Finance* 72.4, pp. 1483–1528.

Perold, A. (2007). “Fundamentally Flawed Indexing”. *Financial Analysts Journal* 63.6, pp. 31–37.

Petajisto, A. (2013). “Active Share and Mutual Fund Performance”. *Financial Analysts Journal* 69.4, pp. 73–93.

Seabold, S. and J. Perktold (2010). “Statsmodels: econometric and statistical modeling with Python”. *9th Python in Science Conference.*

Sebastian, M. and S. Attaluri (2016). “Factor Investing and Adaptive Skill: 10 Observations on Rules-Based Equity Strategies”. *The Journal of Investing* 25.1.

Sharpe, W. (1964). “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk”. *The Journal of Finance* 19.3, pp. 425–442.

– (1966). “Mutual Fund Performance”. *The Journal of Business* 39.S1, p. 119.

Shumway, T. (1997). “The Delisting Bias in CRSP Data”. *The Journal of Finance.*

Thomas, S. H., A. D. Clare, and N. E. Motson (2013). “An Evaluation of Alternative Equity Indices - Part 2: Fundamental Weighting Schemes”. *SSRN Electronic Journal* April.

Treynor, J. (1965). “How to rate management of investment funds”. *Harvard Business Review.*

Treynor, J. and F. Black (1973). “How to Use Security Analysis to Improve Portfolio Selection”. *The Journal of Business* 46.1, p. 66.

Treynor, J. and K. Mazuy (1966). “Can Mutual Funds Outguess the Market?” *Harvard Business Review* 44.4, pp. 131–136.

Wermers, R. (2000). “Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transactions Costs, and Expenses”. *The Journal of Finance* 55.4, pp. 1655–1695.

Wermers, R., H.-L. Chen, and N. Jegadeesh (2000). “The Value of Active Mutual Fund Management: An Examination of the Stockholdings and Trades of Fund Managers”. *SSRN Electronic Journal* 35.3.

Xiong, J. et al. (2010). “The Equal Importance of Asset Allocation and Active Management”. *Financial Analysts Journal* 66.2, pp. 1–9.