General Superposition Strategies and Asset Allocation

by

Kathryn Margaret Kaminski

Submitted to the SLOAN SCHOOL OF MANAGEMENT
in partial fulfillment of the requirements for the degree of
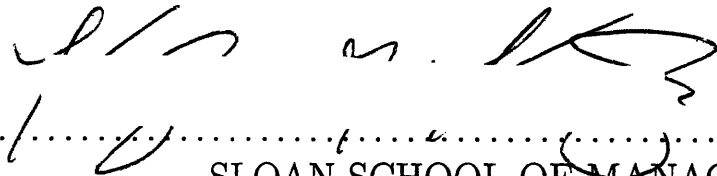
Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2007

Author . . . .
SLOAN SCHOOL OF MANAGEMENT
May 7, 2007

Certified by . . . .
Andrew W. Lo
Harris & Harris Group Professor of Finance
Thesis Supervisor

Accepted by . . . .
Dimitris Bertsimas
Co-Director Operations Research Center

General Superposition Strategies and Asset Allocation

by

Kathryn Margaret Kaminski

Abstract

Investors commonly use stopping rules to help them get in and out of their investment positions. Despite their widespread use and support from behavioral finance, there has been little discussion of their impact on portfolio performance in classic portfolio choice theory. In this thesis, I remedy this situation by discussing the performance impact of stopping rules, highlighting the stop-loss rule.

Stop-loss rules—predetermined policies that reduce a portfolio's exposure after reaching a certain threshold of cumulative losses—are commonly used by retail and institutional investors to manage the risks of their investments, but have also been viewed with some skepticism by critics who question their efficacy. I develop a simple framework for measuring the impact of stop-loss rules on the expected return and volatility of an arbitrary portfolio strategy, and derive conditions under which stop-loss rules add or subtract value to that portfolio strategy. I show that under the Random Walk Hypothesis, simple 0/1 stop-loss rules always decrease a strategy's expected return, but in the presence of momentum, stop-loss rules can add value. To illustrate the practical relevance of this framework, I provide an empirical analysis of a stop-loss policy applied to a buy-and-hold strategy in U.S. equities, where the stop-loss asset is U.S. long-term government bonds. Using monthly returns data from January 1950 to December 2004, I find that certain stop-loss rules add 50 to 100 basis points per month to the buy-and-hold portfolio during stop-out periods. By computing performance measures for several price processes, including a new regime-switching model that implies periodic "flights-to-quality," I provide a possible explanation for our empirical results and connections to the behavioral finance literature.

Consistent with the traditional investor's problem, I discuss a generalization of this approach to general stopping rules, which are superimposed on arbitrary portfolio strategies. I define a stopping utility premium and discuss how uncertainty about the true stochastic process can explain a potential value added or value lost by the use of stopping rules in practice.

Thesis Supervisor: Andrew W. Lo
Title: Harris & Harris Group Professor of Finance

Acknowledgments

I would first like to thank my advisor Andrew Lo for all of his guidance and encouragement during my time at the ORC. I am truly grateful for the opportunity to work with him. I found his dedication to research and his tenacity about new ideas inspiring both professionally and personally. I thank him for encouraging me to learn to think independently and to take risks in my career.

I am also thankful for the continued advice and support of my thesis committee members, Professors Dimitris Bertsimas and Roy Welsch. I was also fortunate to receive excellent advice concerning my transition into finance from Professors Stewart Myers, Dirk Jenter, Gustavo Manso, and Leonid Kogan. I would also like to thank the ORC staff: Andrew, Laura, Paulette, and Veronica.

While at MIT, I was an active member of the Laboratory for Financial Engineering. Through this lab, I met Jasmina Hasanhodzic, Mila Getmansky-Sherman, Mike Epstein, Svetlana Sussman, and Tom Brennan and I am thankful for their friendship and the opportunity to learn more about finance from them. In addition to her support and encouragement as an administrator, Svetlana Sussman is a great friend. I was also actively involved in the Financial Technology Option (FTO). As part of the FTO, I enjoyed the opportunity to work with Debra Luchanin and many excellent students. I also would like to thank Barbara Jensen, from Alpha Simplex, for her support and cheerful spirit.

One of the greatest aspects of my tenure at the ORC has been the incredible people I have known and worked with. I thank Margrèt and Juliane for always supporting and encouraging me - I will miss them dearly. I will also never forget the great times I had with my friends Jose, Carolina, Nico, Mike Sr., Mike Jr., Nebibe, Sebastian, Mohammed, Theo, Ilan, Ruben, Yann, Karima, Nelson, Dan, Charlotte, and Guillaume.

I would also like to thank my non-ORC friends for making my life at MIT one that I will always remember fondly. Even though they were very far away, Lynn, Melanie, and Ann were just a quick phone call away. I also thank Lucile and Rebeca

Contents

List of Figures

List of Tables

Chapter 1

Introduction

As changes in government policy and the reduction of social security shift more and more responsibility of retirement investment decision-making onto the individual investor, effective financial planning and asset allocation has become a critical issue. In parallel, the field of portfolio choice has developed an ever expanding set of solutions to the asset allocation problem. Despite the plethora of solutions, practical limitations relating to the difficulty in forecasting returns and the real costs of active trading suggest a more passive approach to investment, based on low-cost index funds with little emphasis on active trading or frequent rebalancing. Despite the simplicity of this passive approach, human behavior often contradicts this approach leading investors to change their investment positions based on market performance. As a result, in practice, portfolio allocation is dominated by the use of common heuristic strategies and simple rules that overlay underlying investment strategies. In particular, stop-loss rules—predetermined policies that close out long or short positions after reaching a certain threshold of losses—are commonly used by retail and institutional investors to manage the risks of their portfolios. Due to the highly nonlinear nature of these stopping rules, it is difficult to gauge their impact on overall portfolio performance in the classic framework of portfolio choice. In this thesis, I develop a general framework for measuring the efficacy of stop-loss rules and other stopping-based market-timing strategies, which is applicable to general price processes and independent of sampling frequency. I then characterize the link between stop-loss, nonlinear market-timing

strategies, and a stopping premium. This link allows me to provide guidelines for the use of stopping rules in practice, which connects in a very natural way to standard portfolio theory.

This thesis consists of three main sections. In Chapter 2, I discuss practical asset allocation by focusing on four discussion areas: financial planning, traditional portfolio choice theory, behavioral finance, and empirical findings in asset allocation. This section explains some of the key differences between asset allocation in practice and theory, as well as provides motivation for examining stopping rules. Motivated by the portfolio choices of investors and my discussion of asset allocation, I then examine the efficacy of the classic stop-loss rule in Chapter 3. In order to explain "When stop-loss rules stop losses?", I develop a framework for measuring their performance impact on an underlying strategy and derive conditions under which stop-loss rules add or subtract value from an arbitrary portfolio strategy. In particular, I discuss how, under the random walk hypothesis, stop-loss cannot stop losses, whereas, under momentum or regime switching, stop-loss may actually stop losses. To illustrate the practical relevance of my framework, I provide an empirical analysis of a stop-loss policy applied to a buy-and-hold strategy in U.S. equities, where the stop-loss asset is U.S. long-term government bonds. Using monthly returns data from January 1950 to December 2004, I find that certain stop-loss rules add 50 to 100 basis points per month to the buy-and-hold portfolio during stop-out periods. By computing performance measures for several price processes, including a new regime-switching model that implies periodic "flights-to-quality," I provide a possible explanation for the empirical results and connections to the behavioral finance literature. Given the success of my framework for examining the efficacy of stop-loss rules, I extend this framework to more general superposition type stopping rules in Chapter 4. I demonstrate how stopping rules can add or subtract value to arbitrary portfolio strategies. I explain this more specifically by examining scenarios where I relax the assumption that the stochastic process for asset returns is known, and I show how in the classic utility sense stopping rules can add value.

Chapter 2

Asset Allocation

In this chapter, I highlight several important aspects in asset allocation including: financial planning, traditional portfolio choice theory, behavioral finance, and empirical portfolio choice. In Section 2.1, I discuss financial planning for individual investors by examining the current state of the field for long term investment. I explain who provides investment services for households and the specific investment vehicles that are available for the typical household retirement investor. In Section 2.1.4, I discuss how a retirement plan is created and predominant rules of thumb for asset allocation in retirement planning. I then discuss several heuristic strategies, which can be deduced, by examining common retirement investment options as well as general advice in Section 2.1.5. Following my discussion of the practical implementation of financial planning, I then turn to the academic based theory of portfolio choice in Section 2.2. I review the field of portfolio choice and discuss how it has been extended to predictability in asset returns in Section 2.2.1 as well as its connections with behavioral finance in Section 2.2.2. Given the recent emphasis on behavioral finance in portfolio choice, I then turn to the field of behavioral fiance and examine the basic relevant principles for investor decision-making in Section 2.3. After my examination of the practical aspects of portfolio planning in Section 2.1, my discussion of the academic findings relating to portfolio choice in Section 2.2, and a review of important investor decision making principles in Section 2.3, it only seems fit that the final section of this chapter focuses on what household investors' are actually doing in their investment

accounts. I discuss the current findings in empirical portfolio choice allowing me to point out the important characteristics it shares with the principles laid out in the three previous sections.

2.1 Financial Planning

The financial planning field has long been providing advice for both individuals and businesses on how to effectively save and plan their financial future. Most investors are not aware of how to structure their lifestyle in ways that allow them to create effective spending and savings habits that can secure their stable retirement, children's education, or buy their dream house. Individual investors are faced with real world, non-linear and uncertain constraints, which may make taking financial decisions feel daunting and complex. As a result financial planning is both a necessary and complicated experience for investors. Since financial planning is a field which must cater to wide variety of needs and goals for individual investors, the field is governed by simple heuristics, general advice, and simple rules to follow.

Financial planners stress several important but simple goals for investors. Those goals include: creating positive net worth, spending less than you earn, and starting to save as soon as possible. These goals are exactly mirrored by academic work based on positive NPV and the time value of money. The aim of financial planning is to help individuals create good financial habits which are consistent with their financial goals. Since financial planners must attempt to provide a suitable solution financial plan for each individual investor, instead of exact solutions they offer general heuristic advice. In particular, they suggest tax-deferred investing (participation in a retirement plan), diversification, long term contributions, financial protection, and monitoring.

The use of tax exempt vehicles is heavily stressed by financial planners because it provides a simple way for investor to take advantage of "free" opportunities given to them by the government. They suggest making maximum contributions into tax-deferred or tax-exempt funds such as IRAs or employee based defined contribution plans. I provide a summary of these investment options in Section 2.1.2.

Financial planners understand that a tax exempt vehicle is a simple way for an investor to make decent returns without having to make speculative choices. Financial planners also suggest that investors engage in automatic saving strategies. They provide heuristic methods to maintain this behavior which include joining defined contribution plans, saving consistently with every paycheck, diverting earnings directly into savings accounts, and saving part or all of bonuses and raises. This advice is important because it demonstrates how investors can use self-control mechanisms which avoid human behaviors that are not consistent with good financial health. More specifically, by saving systemically, one avoids the potential for putting money into the market only during upward movements in the market. Automatic investing helps the investor "average" into the market.

Financial planners also stress diversification. They remind investors that having some portion of their investment in products with higher risk can help boost returns. This boost in return can help keep individuals up to speed with inflation. As for where to invest, they suggest diversified funds that can provide the portfolio management. There is one point where the investment advice of financial planners is controversial. They make the comment that you can afford more risks if your investment horizon is longer. This comment can sometimes be used to suggest that financial planners think risks decrease in the long run, which is not true. In fact, risks increase with horizon but for most people risk tolerance decreases with horizon.

Financial protection is also important to secure long term investments in the case of an emergency or down-period in the economy. Financial planners suggest that investors maintain an emergency fund of money which is accessible to prevent from early liquidation of long term holdings. Another form of financial protection is insurance. Products such as life insurance, car insurance, and fire insurance can help investors protect their long term investments and help them to deal with financial crises.

Since financial goals, market conditions, and income conditions change with time, financial planners also suggest monitoring and revising financial plans every 2 or 3 years or when there are significant changes in preferences, goals or financial condi-

tions such as changes in the tax code. Monitoring financial performance can help keep investors on track with their goals. The low monitoring frequency of individual investors further suggests that when they invest they should invest in products such as diversified funds which automatically adjust with time or require minimum monitoring frequency.

Financial planners also stress the importance of avoiding common financial traps. The most common of these include high interest debt such as credit cards and loans. Since one of the greatest mistakes investors make is lack of participation in savings plans, they stress that it is never too late to start saving for the future.

2.1.1 Long Term Investment: Current State of the Industry

Over the past several years there have been significant changes in the financial industry which has resulted in changes in the structure of how people invest their money. Banks, Insurance Companies, and Brokerage Firms have become more similar in the services that they offer. As a result, the field of Financial Planning has expanded significantly and grown in importance over the past twenty years.

As a result of financial planning advice, the financial industry has developed products which attempt to cater to the needs of individual investors. Companies like Vanguard®, Fidelity®, TIAA-CREF®, Janus Funds® and American Funds® provide a wide array of funds for individuals with varying preferences. Each of these funds promote their services by suggesting that they provide diversification, portfolio management, lower costs than active investment, experience in the field, size, continuity, and ease for the investor.

Vanguard® suggests there are three important considerations for a potential investor: investment objective, time horizon, and risk tolerance. Common investment objectives are retirement planning, funding college education, estate planning, general investing, and tax planning. Time horizon is important because several funds have time varying allocation between different asset classes. Risk tolerance is considered synonymous with the amount of allocation to stocks versus bond and cash equivalents. The main categories for classifying risk are conservative and growth.

Conservative funds have less exposure to equities and growth funds are more heavily weighted in equities and higher risk products.

Retirement funds or retirement planning funds which vary asset allocation with time horizon are often called lifecycle or lifestyle funds. Each of these funds is essentially a fund of funds which is mixed appropriately for the desired amount of risk tolerance whether it be growth, conservative growth, etc. Common examples of the lifestyle or lifecycle funds are Vanguard LifeStrategy® Funds, and Fidelity Freedom® Funds. A broader review of investment options is presented in the following section.

2.1.2 Retirement and Long Term Investment Options:

When it comes to retirement plans, there are many options in particular: Qualified plans, individual retirement accounts (IRAs), "almost"qualified plans, and Non-Qualified Plans. These plans are generally sweetened with tax-breaks to encourage taxpayers to save for retirement. Because of these "free opportunities", the government has laid out very complicated rules governing how much money can be put into these plans, as well as regulations and penalties for withdrawing monies before retirement, and distribution requirements.

A Qualified Plan is a plan that is qualified to receive certain tax benefits as described by Section 401 of the U.S. Tax Code. It is essentially a forced savings plan established by an employer to benefit its employees. Plans of this type include 401(k), profit sharing plans, Stock Bonus Plans, Money Purchase Pension Plans, ESOPs, Defined Benefit Plans Target Benefit Plans and Self-Employed Plans (Keoghs). The most common of these are the 401(k) Profit Sharing Plans and Defined Benefit Plans.

A 401(k) Profit Sharing Plans or simply 401(k) plans allow the employees to choose to defer some of their salary and employers can also make contributions. Typically, there are "vesting" rules that apply. Employees are always "100%" vested in their deferrals but the employer contributions are subject to various vesting rules. The growth of such plans has shifted more responsibility for retirement savings and investment management onto the individual. In order to provide a diversified set of

bonds and equities, employers typically offer a "menu" of funds that the employee may choose from.

401(k) plans are quite different from the Defined Benefit Plans. In contrast, Defined Benefit Plans place more of the funding burden on the employer rather than the employee. They typically require the employer to make annual contributions and the plan is generally maintained in one large account. The terms of a defined benefit plan always include a promise to pay each participant a specific dollar amount as an annuity beginning at retirement. In order to fund this plan, companies must compute annually the funds necessary to satisfy their retirement liabilities. Recent stock market corrections and the decline in interest rates have resulted in many large companies having "under-funded" defined benefit plans. In addition to having strongly concentrated company risk, Defined Benefit Plans can also be non-transferable. In the modern workplace, investors need more flexibility as well as more diversified risks to protect their investments. As a result, Defined Benefit Plans are becoming less and less common.

Another investment option which employers can offer is a Stock Bonus Plans. Stock Bonus Plans are similar to profit sharing plans with the exception that the employer contributes to the plan in the form of shares of the company stock. It is clear that this can be a problem for employees of companies which have gone under such as MCI® or Enron®. Generally, these plans are available for larger companies who can afford to administer and comply with federal regulations. These plans are offered in hopes of promoting employee loyalty, yet many financial experts would argue that such plans force employees to be insufficiently diversified.

Over the past years, the Federal Government has also provided more options for small businesses, self-employed individuals, and individuals who are not covered by employer sponsored plans. Such opportunities include Individual Retirement Accounts (IRAs), SIMPLE IRA Plans (Savings Incentive Match Plan for Employees of Small Employers), SEPs (Simplified Employee Pensions), and Roth IRAs. As in profit sharing plans, these plans are also subject to regulations governing contributions amounts, withdrawal restrictions, and distribution requirements. The Federal

government eliminated both reporting and annual filings which made these plans more affordable to administer. As a result, these plans are available to individuals and small business through large brokerage firms, banks, insurance companies, and low cost providers such as Vanguard, Fidelity®, and Schwab®. As in a 401(k), the funds available to individual participants are broad and therefore, the individual bears the burden of selecting and managing their investment accounts.

2.1.3 Where to Invest, Who can Invest, and in What?: An Industry Summary

In Table 2.1, I summarize the investment providers and the services they provide to investors. The structure of this table is the result of changes in the tax code, changes in governmental regulations, and changes in reporting rules for investment providers over the past few decades. Unlike in the past, there are now many different sectors or types of providers who are catering to various types of investors. The main types of financial institutions which provide investment services are investment banks, insurance companies, brokerage firms, mutual fund companies, and registered investment advisors. Despite the division in Table 2.1, the lines between types of investment providers has grown increasingly blurred. For example, Fidelity®, Vanguard®, and Schwab® provide mutual funds and brokerage accounts while also offering insurance and annuities. Northwestern Mutual® sells both insurance and mutual funds. Increasingly, insurance agents are becoming certified financial planners in order to be licensed to sell mutual funds. It is clear that the industry is offering a broader spectrum of services, but it seems to be significantly harder for an investor to know where and with whom to invest.

As I explained in Section 2.1, a good financial plan should be highly specialized to suit that individual's needs and goals. Unfortunately, it is the case, in general, that the larger the assets the more individualized the approach. Most investment providers have certified financial planners (CFPs) who create a "financial plan" for investors. This financial plan involves gathering a great deal of information regarding: income,

Financial Institutions	Mutual Funds	Equities/ Fixed Income	Insurance Products	Private Accounts Individualized	Fee Structure
Investment Banks	Yes	Yes	Yes	Yes	Varies
Insurance Companies	Yes		Yes		
Brokerage Firms	Yes	Yes	Yes	Yes	Varies, trading fees % AUM
Mutual Funds	Yes				
Registerd Investment Advisors	Yes	Yes	Yes	Yes	Fees based on % AUM

Table 2.1: A Summary of Financial Investment Providers and Their Products and Services

expenses, debt level, projected growth of income, risks, current assets, retirement assets, estate planning (preparation of wills, etc.), tax issues, financial goals, ability to tolerate risk, investment experience, and commitment to the process. Planners may charge a fee for this service; this fee is sometimes waived if the plan leads to management of assets, sales of insurance products, etc. Since financial advisors are compensated based on the fee structure and company incentives, there are always questions regarding conflict of interest between advisors and clients. In many cases, investors may be unaware of incentive based bias of investment advisors.

2.1.4 Asset Allocation in Retirement Plans

Despite the amazing advances in quantitative methods applied to finance and portfolio theory, the choice of asset allocation for individual investors remains more of an "art" than a science. Since most individuals are unaccustomed to financial markets and financial decision-making, the choice of asset allocation is often decided upon by financial planners based on simple heuristics and general categories of risk tolerance.

For investors retirement planning can be divided into two steps: the financial goal setting and asset allocation. Financial goal setting requires investors to estimate the required funds they will need to support their lifestyle after retirement. An individual must estimate time to retirement, years in retirement, income till retirement, income after retirement, bequest motives, and required rates of return. With all of these

inputs, financial planning tools can help investors estimate how much they should contribute regularly to their 401(K), IRA, and retirement savings accounts. Once investors have a good idea of the rates of return they require, they are still far away from completing their financial plan because they still need to decide what they will invest in.

As Fidelity® boasts on their website, asset allocation is what drives 91% of returns. The choice of asset allocation is a difficult one because financial planners rely on both personal intuition as well as a basic benchmarking scale for defining risk tolerance. The following two heuristics are prevalent in all financial planning strategies provided by investment companies today.

1. Investors can bear larger risks when the horizon is large

2. Asset Allocation to stocks versus bonds is indicative of risk tolerance, i.e. higher tolerance for risk more exposure to stocks versus bonds and cash equivalents.

Investment strategies are benchmarked by how much growth or income they claim to provide. Growth portfolios are heavily weighted in stocks, foreign equities, and higher risk ventures. Income portfolios are heavily weighted in bonds and therefore they tend to fluctuate less in value and provide cash outflows such as coupon payments. Portfolios of both stocks and equities are generally categorized from aggressive to conservative, where aggressive generally denotes more exposure to stocks and conservative denotes more exposure to bonds. Table A.7 demonstrates how Vanguard® explains a heuristic between asset allocation and portfolio objectives for its LifeStrategy® Funds. Fidelity® provides the categories in Table A.8 and their heuristic mapping to asset allocation, they also suggest appropriate time horizons.

To get a better good idea about strategies for asset allocation which are suggested by financial planners, it is beneficial to examine all-in-one mutual fund options for individuals. The most widely cited advantage of an all-in-one mutual fund is that they provide less monitoring frequency and less involvement for individuals. This may be advantageous for individuals with less experience in markets who may be incapable or uninterested in making portfolio decisions over time. All-in-one funds

are also important for this thesis because they provide some transparency in regards to shifts in asset allocation over time. The three main types of all-in-one fund choices are lifecycle funds, balanced allocation funds, and target maturity funds. Each of these fund types are summarized in the following paragraphs.

Target maturity funds are mutual funds which are a basket of index type funds diversified among asset classes. The key characteristic of this type of fund is that the allocation to various asset classes varies as the fund matures. Since investors do not have to monitor or adjust their allocation with time, they do not incur taxable gains, taxes, or trading costs beyond the base mutual fund fees. While being well diversified within asset classes, target maturity funds are not diversified in their mutual fund or index fund providers so they may be some intrinsic company specific risk since all of the indices are offered by the same investment company. In addition, target maturity funds vary substantially in the level of risk which they deem appropriate over different horizons based on the fund provider. For example, one investment company may have a more conservative or more growth oriented tilt to the underlying asset allocation. I will examine how several investment companies implement their time varying allocation to demonstrate how they benchmark their asset allocation over time. Tables A.1, A.2, A.3, A.4, and A.5 list the time dependent allocation specification for target maturity funds for Fidelity®, Vanguard®, Putnam Investments®, American Century Investments®, and T. Rowe Price® respectively.

Lifecycle or lifestyle funds are similar to target maturity funds in that they are well diversified among asset classes yet they maintain a relatively static allocation strategy over time. Lifecycle funds are differentiated by general categories of risk tolerance as defined by growth, conservative growth, income, etc. Tables A.7, A.8, A.9, A.10, and A.11 list the allocation for lifecycle type funds for Vanguard®, Fidelity®, T. Rowe Price®, Schwab®, and American Century Investments® respectively.

Balanced funds are similar to lifecycle funds except that they are funds which invest in securities directly in order to track index performance. They are generally focused on their relative performance with respect to a benchmark index or basket of indices. Allocation specifications for balanced funds for Fidelity®, T. Rowe Price®,

and Schwab® are also listed in Tables A.8, A.9, and A.10 respectively.

2.1.5 Financial Planning Heuristics for Asset Allocation

Given my general discussion of heuristics for asset allocation in the previous section, in this section I outline in further detail several common asset allocation strategies.

Buy-and-Hold

By far the most common and the most simple strategy is the buy-and-hold strategy. In a buy-and-hold strategy, the investor simply decides the initial amount of stock and bonds to purchase and he or she does not alter the allocation throughout the course of the investment. The buy-and-hold strategy is motivated by the assumption that the initial allocation will be optimal if it is held over a long enough period of time. Financial Planners advise investors to avoid speculation about market performance since market moves may be temporary. These temporary moves may cause investors to be pushed into action by simple noise effects. A buy-and-hold portfolio may make more sense for an investor who is interested in stock only or bond only portfolios because it does not require intermediate rebalancing. Common examples of buy-and-hold portfolios are those of Stock funds, Index Funds, Bond funds, etc. Mathematically, the evolution of the portfolio weights w_t at time t for the risky asset S_t can be described by the following equation where B_t denotes the value of bonds at time t, and n^S and n^B denotes the initial number of stocks and bonds purchased.

$$w_t = \frac{n^S S_t}{n^S S_t + n^B B_t}$$

Value Weighted with periodic adjustment

Due to the increased duration of retirement funding as well as the long lived retirement income plans, most financial planners would suggest that horizon effects have diminished. This suggests that investors generally choose a level of risk tolerance with which they are comfortable. Once used to this level of tolerance, they will remain

comfortable with that level unless they are extremely close to the end of their lives or there is a large market event which can change general investor sentiment. This suggests that a risk tolerance appropriate asset allocation mix would be an appropriate strategy. As a result, many investment management companies offer a wide selection of balanced funds or Lifecycle funds with benchmark allocation such as 60 40. To avoid excessive trading costs and to attempt to profit from some short term market timing, most fund managers have either implicit or explicit (as defined in the investment prospectus of the fund) bounds or limits on their exposure to various asset classes. In theory, these bounds can help managers to sell when stocks are high and buy when stocks are low. Classic examples of this type of strategies are the so called lifecycle funds as discussed in the previous section. Table A.11, American Century Investments® One Choice Portfolios provides an example of explicit bounds on the allocation to asset classes. Under a value-weighted strategy, the evolution of allocation weights w_t to the risky asset S_t can be described in the following equation where w_u and w_l are the upper and lower bounds on allocation to the risky asset respectively. n_t^S and n_t^B denote the number of stocks (or shares of an index) and bonds in the portfolio at time t. S_t and B_t are the prices of the risky and riskfree assets at time t.

$$
w_t = \begin{cases} w_u & \frac{S_t n_t^s}{S_t n_t^s + B_t n_t^b} \geq w_u \\ \frac{S_t n_t^s}{S_t n_t^s + B_t n_t^b} & w_l < \frac{S_t n_t^s}{S_t n_t^s + B_t n_t^b} < w_u \\ w_l & \frac{S_t n_t^s}{S_t n_t^s + B_t n_t^b} \leq w_l \end{cases}
$$

Lifestyle Planning

Although the link between risk tolerance and age is actively debated, financial planners still argue that risk tolerance should decrease with age. For example, with a shorter horizon, they argue that you have less opportunities to wait out unfavorable market conditions. Because of age effects, several financial companies have marketed target maturity funds which become progressively more conservative as the lifespan of the fund decreases. A simple heuristic consistent with a target maturity fund is

the classic age heuristic. Simply put the age heuristic is to allocate 100 minus your age to stocks. The evolution of the portfolio weights to stocks w_t can be described for the age heuristic in the following equation with a_t denoting the age of an individual at time t.

$$w_t = \frac{a_t}{100} - 1 \quad Age\ Heuristic$$

Most target maturity funds divide assets into three basic asset classes: stocks, bonds, and cash equivalents. If I define w_t^S, w_t^B, and w_t^C as the allocation to stocks, bonds, and cash equivalents, I can describe their evolution in over time in the following equation where T is the total retirement period, $f_S(\cdot), f_B(\cdot)$ and $f_C(\cdot)$ are functions of the remaining time to maturity $t - T$ and (w_0^S, w_0^B, w_0^C) are the initial allocations to stocks, bonds, and cash equivalents respectively.

$$
\begin{aligned}
w_t^S &= w_0^S - f_S(T - t) \\
w_t^B &= w_0^B - f_B(T - t) \\
w_t^C &= w_0^C - f_C(T - t) \\
where\ 1 &= f_S(T - t) + f_B(T - t) + f_C(T - t)
\end{aligned}
$$

Tables A.1, A.2, A.3, A.4, and A.5 in the appendix demonstrate how $f_S(T-t), f_B(T-t)$ and $f_C(T - t)$ evolve as a function of time to maturity $T - t$. It is clear from their prospectuses that stock allocation decreases with time, bond allocation increases with time, and that short term products are introduced only for time periods 15 or 10 years or less. It is also important to note that for larger horizons (25 years or greater) there are minimal changes in allocation and most funds have relatively aggressive growth strategies. Given the rough description of how allocation varies over time, the time varying allocation can be simplified into a set of linear constants α and β with initial

and final allocation limits w_1 and w_2.

$$
w_t^S = \begin{cases} w_1^S & T - t \geq t_1^S \\ \alpha^S(T - t) + \beta^S & t_1^S < T - t < t_2^S \\ w_2^S & T - t < t_2^S \end{cases}
$$

$$
w_t^B = \begin{cases} w_1^B & T - t \geq t_1^B \\ \alpha^B(T - t) + \beta^B & t_1^B < T - t < t_2^B \\ w_2^B & T - t < t_2^B \end{cases}
$$

$$
w_t^C = \begin{cases} w_1^C & T - t \geq t_1^C \\ \alpha^C(T - t) + \beta^C & t_1^C < T - t < t_2^C \\ w_2^C & T - t < t_2^C \end{cases}
$$

$$
where \ 1 = w_t^S + w_t^B + w_t^C
$$

A summary of linearized weights for several investment companies is detailed in Table A.6.

Constant Proportion Portfolio Insurance (CPPI)

A CPPI strategy buys shares as they rise and sells shares as they fall. The investor selects a floor below which the portfolio value is not allowed to fall. The floor increases with the rate of return of cash. The cushion is defined as the difference between the portfolio value and the floor. A CPPI decision rule is to keep the exposure to risky assets at a constant multiple of the cushion. Letting w_t be the weight in the risky asset at time t, f be the floor value, and m be the multiplier, the weight in the risky asset w_t can be described by the following equation.

$$
w_t = \frac{(S_t n_t^s + B_t n_t^b - f e^{rt}) * m}{S_t n_t^s + B_t n_t^b}
$$

Market Timing

Since overall asset allocation has been cited as the driving factor for performance, it is no surprise that some level of market timing could incur improvements in performance.

For asset allocation, the relative performance of bonds to stocks is a key factor for adjusting allocation. This suggests a simple heuristic that stipulates that when debt provides low returns, allocation should be slanted more into equities. When the opportunities in bonds are attractive, the allocation should be slanted more into bonds. A well-known example of a time-varying allocation heuristic between stocks and bonds is the Fed Model. The Fed Model compares the yield on a 10 year treasury note and the earnings on the S&P 500 Index. In a market timing allocation strategy, the portfolio weight is a function $f(S_t, B_t, y_t)$ of the underlying asset classes stocks S_t and bonds B_t and a set of state variables y_t at time t. Similar to a value weighted strategy, this strategy may also dictate allocation decisions based on simple bounds on the state variable. Letting y_u and y_l be the upper and lower bounds on the state variable y_t, I can describe evolution of portfolio weights in the following equation.

$$w_t = f(S_t, B_t, y_t)$$

$$w_t = \begin{cases} w_1 & y_t \geq y_u \\ \frac{S_t n_t^s}{S_t n_t^s + B_t n_t^b} & y_l < y_t < y_u \\ w_2 & y_t \leq y_l \end{cases}$$

2.2 Portfolio Choice Theory

The field of portfolio choice and asset allocation began with the work of Markowitz (1952) on mean variance analysis. Mean variance analysis provided the first analytic framework for examining portfolio choices in an analytic manner for single period investment horizons. Mean variance analysis suggested that all investors hold the same portfolio and that the proportion held in risky assets would depend on the risk aversion levels of different investors. Mean variance analysis paved the way for the work of Tobin (1958) on mutual fund separating theorems. Despite the success of mean variance analysis, it could not provide insight for multi-period investment horizons or higher order characteristics of asset returns. As a result, both Samuelson (1969) in the discrete case and Merton (1969) in the continuous case provided multi-period methods for solving the portfolio choice problem using dynamic programming

33

techniques. As research in behavioral finance and predictability in asset returns have evolved, the portfolio optimization problem has been extended accordingly to include non-myopic utility functions, various types of predictability in asset returns, as well as stochastic opportunity sets. I outline the effect of predictability and time-varying opportunity sets on portfolio choice problems in Section 2.2.1. In Section 2.2.2, I provide an overview of recent literature concerning behavioral aspects of portfolio choice.

2.2.1 Portfolio Choice and Predictability

In recent years, as a result of the publication of widespread evidence of predictability in markets (see Brennan and Xia (2001), Xia (2001)), there has been an increased interest in the effect of predictability on portfolio optimization. Not surprisingly predictability has been found to play a significant role in the choice of optimal policies for portfolio construction. The problem of predictability has been addressed both numerically and explicitly under various formulations of the investor's problem. I outline the literature for explicit, approximately explicit, and numerical solutions to the portfolio optimization problem.

There are three well known closed-form solutions to the portfolio choice problem with time-varying opportunity sets: Kim and Omberg (1996), Wachter (2002), and Liu (1999). Using an Ornstein-Uhlenbeck process for the risk premia which is not perfectly correlated with the underlying return, Kim and Omberg (1996) solve the problem explicitly, in the case of terminal wealth with a HARA utility function. Using their explicit solution, they show that investors will hedge against changes in the risk premia and the size of the hedge will depend on the parameters of the underlying dynamics. They also conclude that risky allocation does not necessarily increase with the horizon unless the risk premia is positive. Following the work by Kim and Omberg (1996), Liu (1999) proposed a new formulation of the investor's problem where the risky asset and the instantaneous variance both follow a diffusion process. He solves the problem explicitly and uses the solution to compare static and dynamic portfolio choices. In contrast to standard mean variance analysis, he suggests that

risky allocation may not be a good proxy for risk aversion, and that volatility may not always deter a risk averse investor from risky assets. Using the assumption that markets are complete, Wachter (2002) provides a solution to the investor's problem with utility over consumption in the case of mean reversion in the risk premia. Her solution is important because it provides a method for resolving the issue of terminal wealth versus intermediate consumption. It also provides some key economic insight parallel to traditional theory in fixed income notably coupon bonds, duration, etc. She shows that the optimal portfolio allocation is analogous to a weighted average, similar to duration in coupon bonds. This interpretation is useful because it can explain the sign of the hedging demand, the shortfalls of log-linear approximation approaches commonly used for approximating the budget constraint, as in Campbell and Viceria (1999), as well as convergence results for long horizons.

To circumvent analytical intractability issues in portfolio optimization, several other approximations have been applied to provide approximately explicit solutions. Campbell and Viceria (1999) assume an infinitely long lived investor and provide an analysis of predictability by using a log-linearization of the first order conditions and budget constraints. Other papers have used an expansion of the value function for power utility functions (see Das and Sundaram (2000) and Kogan and Uppal (2002)).

Approximation techniques have been useful in specific cases of the portfolio optimization problem, but most commonly the problem has been solved by exact or approximate dynamic programming techniques. Popular examples include Brennan, Schwartz, and Lagnado (1997), Balduzzi and Lynch (1999), Barberis (2000), Xia (2001), and Brandt (1999). Brennan, Schwartz, and Lagnado (1997) use a joint markov process for three state variables: long-term bonds, short-term bonds, and the dividend yield. Their analysis provides encouraging evidence that predictability has a significant impact on portfolio performance and highlights the importance of time horizon. Balduzzi and Lynch (1999) examine the affect of transaction costs on allocation decisions. Barberis (2000) incorporates parameter uncertainty. Xia (2001) demonstrates how predictability and learning have a substantial affect on portfolio selection. Her analysis suggests that asset allocation should take into account pre-

dictability as well as be dynamically updated as more information about the predictive variables becomes available. Brandt (1999) presents a solution that is robust to distributional assumptions following a non-parametric approach. Ait-Sahalia and Brandt (2001) extend the approach in Brandt (1999) over a larger set of utility functions and examine predictive variable choice without distributional assumptions. Brandt, Goyal, Santa-Clara, and Stroud (2005) provide a simulation based method for discrete problems with state dependence using non-parametric regression methods over simulated paths. Using a log-linearization, Campbell, Chan, and Viceria (2003) solve the problem with many state variables.

Despite the various methodologies for solving the portfolio optimization problem, whether the solution is analytical, approximate, or numerical, it is clear that forces such as predictability, model specification, and model uncertainty can have a significant impact on portfolio performance when compared with the myopic solutions derived by Merton (1969) and Samuelson (1969). In summary, the solution to the portfolio choice problem can still be described in terms of two investments: a myopic investment and hedging terms, as suggested by Merton (1971). However the most recent literature establishes the fact that in many scenarios these hedging terms, whether they are due to changes in the opportunity set or to model uncertainty, can dominate the effects of myopic investment choices. These results suggest that, in many cases, there is value-added to strategies which involve dynamic allocation and market timing effects. Given the clear dependence on assumptions about stochastic processes of each optimal portfolio policy and the reality that the true future evolution of asset returns is both unknown and uncertain, the difficult question still remains for an investor: where and when is a dynamic non-myopic strategy applicable?

2.2.2 Portfolio Choice and Behavioral Finance

Despite the plethora of solutions to the portfolio optimization problem, few deviate from the standard set of utility functions derived by expected utility theory. Given that behavioral finance experts would agree that investors do, in fact, violate several axioms of expected utility, it is clear that there are strong limitations

on the power of these solutions to explain or provide financial advice consistent with actual investor behavior. Behavioral finance theorists have developed a set of alternative utility functions, commonly called non-expected utility, which are more consistent with behavioral choices. Common frameworks for modeling non-expected utility preferences include loss aversion and prospect theory (Kahneman and Tversky (1979,1992)), disappointment aversion (Bell (1985), Gul (1991)) ambiguity aversion (Gilboa and Schmeidler (1989)), and regret theory (Bell (1982), Loomes and Sugden (1982)). As a result of this work, several other authors have examined portfolio optimization using behavioral finance to provide insight for some common financial anomalies and the current state of institutional portfolio construction. Some of these recent applications include loss aversion and prospect theory (Benartzi and Thaler (2001), Berkelaar and Kouwenberg (2000), Berkelaar, Kouwenberg, and Post (2003), Ait-Sahalia and Brandt (2001)), ambiguity aversion (Ait-Sahalia and Brandt (2001)), downside-risk (Berkelaar and Kouwenberg, 2000), and disappointment aversion (Ang, Bekaert, and Liu (2005)). In the following sections, I outline several behavioral based utility functions and review research findings.

Prospect Theory and Loss Aversion

Prospect theory, a framework for modeling behavioral preferences was first laid out by Kahneman and Tversky (1979). Loss aversion is a specific case of prospect theory. Loss aversion captures the fact that people do not weigh losses and gains equally, but in fact they are more interested in relative performance. The notion of relative performance implies that investors are less concerned with absolute wealth. Under loss aversion, individuals weigh the probabilities of events objectively; however, at the same time, they often weigh outcomes unobjectively based on the principles of standard expected utility theory. The effect of loss aversion on portfolio choice has been examined by Benartzi and Thaler (1995), Berkelaar and Kouwenberg (2000), Berkelaar, Kouwenberg, and Post (2005), and Ait-Sahalia and Brandt (2001).

Berkelaar, Kouwenberg, and Post (2005) evaluate the portfolio strategy for both a loss averse utility function using a kinked power function and the general prospect

theory utility function. They find that loss averse investors follow portfolio insurance strategies with risky allocation increasing with horizon. They demonstrate how the convexity of a prospect theory utility function can result in large gambling effects. In an earlier working paper by Berkelaar and Kouwenberg (2000), the authors examine how prospect theory and loss aversion preferences are effected by higher order moments of returns (i.e. skewness and kurtosis). In the case of prospect theory preferences, they find that higher order moments can create substantial break even effects. Ait-Sahalia and Brandt (2001) examine partial predictability, by comparing optimal portfolios from expected utility theory (mean variance and CRRA) and non-expected utility (ambiguity aversion and prospect theory). Consistent with financial planning advice, they suggest allocation dependent on conditional state variables using a time-varying combination of indices. In the case of loss aversion and prospect theory, they find horizon effects that can cause investors to avoid risky allocation for shorter horizons. Regardless of preferences, they find that optimal portfolios include significant market timing. In particular, horizon effects and market timing are the most pronounced for prospect theory investors. They also make another poignant remark about return predictability, warning that it is small in recent years and extremely noisy. This remark highlights the fact that although we have analytically tractable results such as those proposed by Wachter (2002) and Kim and Omberg (1996), the implementation of an optimal strategy, based on predictability, could be extremely difficult, and may result in poor performance. They suggest that their approach using the Euler equations to estimate portfolio weights directly, is more robust to model specification. By avoiding distributional assumptions, they focus on the objective function and allow returns to vary nonlinearly with predictive variables over time.

Ambiguity Aversion

Ambiguity aversion is based on the idea that investors may not be able to assign probabilities to future returns. Ait-Sahalia and Brandt (2001) find that an increase in ambiguity aversion is parallel to an increase in risk aversion. They also find that

increased ambiguity may cause some investors to avoid taking particular positions, for example, in bonds. In the case of risky bonds, the return may not be sufficient to outweigh the ambiguity of investing in them.

Disappointment Aversion

As an alternative to loss aversion, some authors have examined the use of disappointment aversion following the framework set out by Gul (1991). As in loss aversion, bad events are weighted more heavily than more good events. The key difference is that the reference point for disappointment aversion, is simply the certainty equivalent. Disappointment aversion has characteristics similar to expected utility theory and thus is more tractable analytically than many other behavioral models of preferences. In addition, loss aversion introduces convexity into the objective function causing risk seeking, gambling type effects, whereas disappointment aversion maintains concavity. Ang, Bekaert, and Liu (2005) use disappointment aversion to explain non participation in the stock market, as well as, dislike for negative skewness.

2.3 Behavioral Finance

Scholars of behavioral finance and financial psychologists have documented a plethora of long-standing and persistent behavioral effects in all types of investors. For simplicity, these effects can be divided into two main categories: cognitive and emotional. Cognitive effects refer to human cognitive decisions based natural brain processing, which are inconsistent with so called "rational" decision-making. On the other hand, emotional effects are decisions which are linked to strong emotional responses. Both cognitive and emotional effects can be divided into collective or group mentality based decision-making, and individual effects. In fact, individual and collective effects can be quite different in nature. Common individual cognitive biases include anchoring, attention bias, framing, habits, and home bias. Common collective cognitive bias can include consensus, common beliefs, and social learning, etc. Common individual emotional biases include denial, greed, regret aversion, overconfidence, pride, etc.

Cognitive Bias

Individual
Anchoring, Attention Bias, Attribution, beliefs, cognitive overcharge, cognitive dissonance, framing, heuristics,small numbers representativeness, mental accounts, habits, hindsight bias, home bias

Collective

Cascades, common beliefs, consensus, manipulation, memes, mimicry, paradigms, percolation, positive feedback, social learning

Emotional Bias

Individual

Addition, endowment effect, denial, greed, fear, hope, loss/ regret aversion, magical thinking, optimistic bias, overconfidence, pride, status quo bias

Collective

Conformity, crowd hysteria, epidemics, fads, groupthink, herding, peer pressure

Table 2.2: A Summary of Cognitive and Emotional Behavioral Effects: Collective and Individual Biases (Source: Behavioral Finance Definitions http://perso.orange.fr/pgreenfinch/bfdef.htm)

Common collective emotional biases include epidemics, fads, herding, peer pressure, etc. A summary of common cognitive and emotional behavioral effects is presented in Table 2.2. Although it is clear that both cognitive and emotional effects are persistent in investors, in the past few years researchers started to piece together both how they affect individual choices as well as the market as a whole. The task of understanding finance from a behaviorial perspective is difficult, because there are many types of behavioral effects and they vary greatly with the state of the market, by the individual, as well as by their interpretation.

There are two main focuses in behavioral finance, the first being explaining common market anomalies and the second being explaining investor decision-making and

choices. In this thesis, I focus more on the second because it is more relevant to financial planning and individual investment choices. In Section 2.3.2, I will discuss how individual preferences have been modeled and provide motivation as well as a literature summary.

2.3.1 Behavioral Effects

Before attempting to apply behavioral finance to areas such as financial planning and portfolio choice, it is important to understand the basic effects which have been documented in order to motivate the study of their application. Due to the extensive nature of this material, I provide a brief summary of behavioral finance. For further reference, there are several authors which provide a more extensive summary of behavioral finance including Shefrin (2005) and Barberis and Thaler (2003).

2.3.2 Modeling Preferences and Decision Making

The resounding support for the persistent evidence of behavioral biases produced an excellent question for researchers: "How can behavioral biases be characterized in terms of preferences?". This question was first examined in the founding work by Kahneman and Tversky (1979) who defined an alternative method for explaining investor preferences, namely, prospect theory. There have been many extensions to their work to explain a wider base of behavioral phenomena. In addition to prospect theory, other theories and alternatives to expected utility theory include disappointment aversion (Gul(1991)), regret theory (Bell (1982), Loomes and Sugden 1982)), and rank dependent utility (Quiggin (1982), Segal (1987,1989), and Yaari (1987)). In the following sections, I summarize and formalize, where applicable, several of these behavioral based approaches to modeling investor preferences.

Prospect Theory

Prospect Theory, as first proposed by Kahneman and Tversky (1979, 1992), is an alternative method for modeling preferences of investors with two key characteristics.

First individuals are less concerned with absolute wealth but instead they derive utility in terms of gains and losses with respect to a particular reference point. Second, instead of having concave utility functions individuals have S-shaped utility functions with risk averse behavior over gains and risk seeking behavior over losses. As compared to expected utility theory, prospect theory can explain the purchasing of insurance and lottery tickets. Another key feature of prospect theory is that individuals do not always properly weigh probabilities. In fact they often overweigh low probability events and underweigh high probability events. This feature of investor behavior can cause investors to combine very safe and very risky assets resulting in a lack of diversification (Barber and Odean (2000), Benartzi and Thaler (2001), Polkovnichenko (2002)).

Prospect theory's key contribution is the introduction of the concept of framing. Because investors frame their investment choices and engage in mental accounting, they consistently make decisions which are inconsistent with expected utility theory. Thaler (1985, 1999) discusses how mental accounting can cause investors to separate individual investments, and thus they can make decisions which as an aggregate do not satisfy principles of standard expected utility theory.

Inorder to clarify how prospect theory is implemented, I will provide the formulation suggested by Kahneman and Tversky (1979,1992). The first simplification of the problem is to examine a gamble with at most two non-zero outcomes (see also Barberis and Thaler (2003)). They define a gamble as a group of four quantities $(x, p; y, q)$ where an individual gets x with probability p and gets y with probability q with $x \leq 0 \leq y$ or $y \leq 0 \leq x$. The value of the gamble is defined by the following expression.

$$V(x, p; y, q) \;\; = \;\; \pi(p)v(x) + \pi(q)v(y)$$

According to Kahneman and Tversky (1979), the function $v(\cdot)$ is S-shaped and can

be formulated using two functions with $f(\cdot)$ concave and $g(\cdot)$ convex.

$$v(x) = \begin{cases} f(x) & x > 0 \\ g(x) & x < 0 \\ 0 & o.w. \end{cases}$$

The formulation of $v(\cdot)$ allows for risk aversion over gains and risk seeking behavior, or gambling effects over losses consistent with human behavior. When the convexity over losses is more pronounced than the concavity over gains, an individual is said exhibit loss aversion. The second function in the value function, $\pi(\cdot)$, accounts for the weighing of probabilities allowing them to be non-linear. Individuals tend to overweigh small probability events, hence they buy insurance and often pay too much for it. This overweighing of small probability events can be achieved by setting $\pi(p) > p$ when p is small. In addition to small probability events, people place more weight on events that are relatively certain.

Kahneman and Tversky (1979) extended their framework to multiple gambles, by allowing for a set of gambles each with outcome x_i and probability p_i where the value function can be written as follows.

$$V(x,p) = \sum_i \pi_i v(x_i)$$

where

$$v(x) = \begin{cases} x^\alpha & x \geq 0 \\ -\lambda(-x)^\alpha & x < 0 \end{cases}$$

with

$$\pi_i = w(P_i) - w(P_i^*)$$
$$w(P) = \frac{P^\gamma}{(P^\gamma + (1-P)^\gamma)^{1/\gamma}}$$

P_i is the probability the gamble with yield an outcome at least as good as x_i and P_i^*

is the probability the outcome will be strictly greater than x_i. Using experimental evidence Kahneman and Tversky (1979) estimate the parameters under prospect theory preferences to be $\lambda = 2.25, \alpha = 0.88$, and $\gamma = 0.65$. Here λ is the coefficient which determines the level of loss aversion.

Disappointment Aversion

Disappointment Aversion is defined by

$$U(\mu_W) \;=\; \frac{1}{K}\Big(\int\limits_{-\infty}^{\mu_W} U(W)dF(W) + A \int\limits_{\mu_W}^{\infty} U(W)dF(W) \Big)$$

where $U(W)$ is the felicity function, for example Power Utility, and A is the coefficient of disappointment aversion, $F(W)$ is the cumulative distribution for wealth with μ_W as the certainty equivalent for wealth W where K is a scalar given by the following equation.

$$K \;=\; P(W \leq \mu_W) + AP(W > \mu_W)$$

For values of A, where $0 \leq A < 1$, outcomes below the certainty equivalent are downweighted.

Regret Theory and Cognitive dissonance

The concept of regret is one that all individuals are familiar with, but its potential implications on investor behavior was first examined by Bell (1982) and Loomes and Sugden (1982). Both Bell and Loomes and Sugden argued that individuals maximize a modified utility function as a result of regret and other psychological effects. Consistent with Bell's model for regret, Shefrin and Statman (1985) explain that regret can motivate individuals to defer the selling of losing positions and accelerate the selling of winning positions. Similar to regret, cognitive dissonance is the mental dilemma that individuals face when they face the possibility that they may be wrong. Because of regret and or cognitive dissonance, individuals may engage in transactions which

are not consistent with expected utility to reduce or avoid regret. Thaler (1980) suggests that regret is stronger than pride and thus individuals may prefer inaction over action when faced with losses.

Mental Accounting

Thaler (1985, 1999) explains that people use implicit accounting systems. In their investments, they mentally code their gains and losses according to prospect theory. People frame the outcomes in a way that makes them the happiest. Different individuals have varying preferences for the organizational structure of their "mental accounts." As a result, contrary to standard economic assumptions about wealth, wealth is not fungible. In accordance with prospect theory, the difference between losses and gains as well as their relative strengths when they are combined will cause individuals to aggregate outcomes in some cases, and segregate in others. Thaler (1985) defines several cases: multiple gains, multiple losses, mixed gains, and mixed losses. For multiple gains, people prefer segregation because of concavity, where as for multiple losses they prefer integration. For mixed gains, people prefer to cancel losses with large gains whereas for mixed losses people prefer to segregate large losses from small gains. Increases in gains are segregated while increases in losses tend to be integrated. A decrease in a gain is integrated, while an decrease in a loss is segregated. Since individuals tend to examine accounts with a loss separately, the decision to close an account at a loss may be extremely difficult for an investor. Johnson and Thaler (1990) examine the "break even" effect and show that in the case of closing out a loss, an investor will engage in risk seeking behavior and be reluctant to realize the loss.

Self-Control

In the context of financial planning, Shefrin and Thaler (1981) discuss the concept of self control in individuals. They model an investor as a combination of planners and doers. They show that individuals rationally choose to impose constraints on their behavior especially when the benefits and costs are uncertain and occur at different

times. Since most people commonly use rules of thumb when they invest, it is clear that these rules can not be characterized by first order conditions consistent with the standard economic framework. Shefrin and Thaler's findings are consistent with Stigler (1966), who asserts that people find utility in protecting themselves from a future lack of will-power. Given the framework of self control, simple orders like stop-loss orders can be seen as a rule which is imposed to protect an investor from a future inability to sell in the event of a loss.

Selling vs. Buying

Although most theoretical models view buying and selling as similar transactions, in fact, psychologically they are very different. Odean (1999) highlights these differences and explains why behavior the buy and sell side are characteristically different. He explains that on the buy side, with a large pool of potential choices, people engage in attention focusing and they end up buying assets who have gone up in the last two years, seemingly believing that the trend will continue. On the other hand when selling securities, investors tend to sell after a long period of upward trend as well but the investments they sell seem to outperform those they bought after they sell them. Odean (1999) highlights these differences by pointing out that investors are interested in past performance when they sell and they are strongly motivated by their aversion to taking a loss on a sale causing them to prefer to sell the winners in their portfolio.

Disposition Effect

The disposition effect, the disposition to sell winners too soon and hold losers too long, has been examined by Shefrin and Statman (1985). Shefrin and Statman (1985) suggest that the disposition effect can be attributed to loss aversion, mental accounting, regret aversion, lack of self control, as well as tax considerations. They exert that tax implications alone are not sufficient to explain the effect. They suggest that lack of self control against regret aversion could explain the use of controlled behavior choices such as a stop-loss order. They also reason that tax loss selling contrary to the work of Constantinides (1983,1984), is a form of self control because December

46

can be perceived as a tax planning deadline. They also explain how the combination of multiple behavioral attributes of investors could produce such effects. First, Loss aversion where investors are reluctant to realize losses will cause them to engage in risk seeking behavior over losses and to be quicker to realize winners than losers. Investors frame their invests over gains and losses with respect to a reference point such as the purchase point. Secondly, investors engage in mental accounting so they do not aggregate their investment decisions but instead with each purchase a new mental account is opened (Thaler 1985, 1999). In fact, Johnson and Thaler (1990) explain that investors encounter considerable difficulty closing an account at a loss. Third, Shefrin and Statman (1985) discuss the psychological implications of a buying and selling. They suggest that investors are motivated by pride but anxious about regret over decisions they may have made which may not have been correct. They also note that it has been shown by Thaler (1980) that regret is more powerful than pride causing investors to prefer inaction to action when faced with regret. Fourthly, Shefrin and Statman emphasize the human desire for self control. They cite key examples in the behavior of individuals and in the behavior of professionals who follow iron-clad, "cut your losses at x", rules.

As a result of the groundbreaking work of Shefrin and Statman, many authors began to examine the existence of disposition effects in various investor accounts and on market prices. More specifically, several authors have examined the existence of the disposition effect in investor accounts including individual investment accounts (Odean (1998, 1999)), in Finland (Grinblatt and Keloharju (2001)), professional investors (Shapira and Venezia (2001)), and real estate markets (Genesove and Mayer (2001)). Using data from discount brokerage accounts, Odean (1998) showed empirically that individual investors realize winners much more often than losers, except in December due to tax implications. In a following paper, Odean (1999) described adverse trading habits of individual investors. He found that the stocks investors sell, outperform the ones they buy. He discussed several regularities in the sample. First, investors buy securities with greater price changes over the last two years than the ones they sell. Second, they buy a similar number of winners and losers but they

sell more of the winners. Lastly, they sell securities that have risen sharply in the weeks prior to sale. Odean suggests that these patterns can be explained, in part, by the psychological differences between buying and selling. He suggests that when investors buy securities they engage in attention focusing, because they have so many securities to choose from. As a result, they are attracted to those securities with abnormal performance. In addition, when investors are selling a security, because of the lack of short selling, they are limited to the small set of securities they own. The adverse selling behavior can be attributed to the reluctance to short sell and the disposition effect. Because of behavioral attributes, the decision to sell involves both past and future performance, whereas in a buy situation an investor only needs to form future expectations. What we can gain from Odean (1999) is that the act of selling can be exceedingly complex for an investor, because it strongly affected by behavioral biases. On the other hand, he suggests that the main behavioral bias on the buy side is attention focusing. The divide between buying and selling provides a clear motivation for self control tactics. Self-control mechanisms help an investor predetermine action on the sell side, when an investor is most influenced by regret, loss aversion, and risk-seeking behavioral tendencies.

Based on empirical evidence it is clear that the disposition effect exists, and thus it must have implications on market prices. In Odean's work on trading behavior, he suggests that the trends in trading behavior suggest that individuals do have information, but they are not using it correctly. In addition to Odean's work on the disposition effect several other authors have examined the its potential implications. Jegadeesh and Titman (1993) show that investors buy at the top of momentum cycles where the trends tend to reverse within a year. They show that momentum based strategies can generate profits of 1% per month. Nofsinger and Sias (1999) find price reversals for those securities with a high percentage of individual ownership. Ranguerova (2001) showed that disposition effects are increasing in market capitalization. Grinblatt and Han (2004) examine a model of the reference point distribution to explain that these behavioral forces create a spread between the fundamental value and the stock price.

2.4 Empirical Portfolio Choice

Despite the many possible solutions to the portfolio choice problem, industry recommendations and educational resources for financial planning suggest a more passive approach to investment, based on low-cost index funds with little focus on active trading or frequent rebalancing. This "buy-and-hold" philosophy is supported by the prevalence of common investment strategies and the recent popularity of all-in-one asset allocation and target retirement funds.

Although a passive approach to investment seems over simplified in the field of portfolio choice, empirical studies have demonstrated how human behavior is inconsistent with this approach. In particular, household investment behavior is fraught with drastically simpler issues. Ameriks and Zeldes (2004) observe that:

> ... a great deal of observed variation in portfolio behavior may be explained by the outcome of a few significant decisions that individuals make infrequently, rather than by marginal adjustments continuously.

Moreover, other documented empirical characteristics of investor behavior include: non-participation (Calvet, Campbell, and Sodini 2006), under-diversification (Calvet, Campbell, and Sodini 2006); limited monitoring frequency and trading (Ameriks and Zeldes 2004, Agnew, Balduzzi, and Sunden 2003), survival-based selling decisions or a "flight to safety" (Agnew 2003), an inability to hedge risks (Massa and Simonov, 2004), and concentration in simple strategies through mutual-fund investments (Calvet, Campbell and Sodini 2006). Variations in investment policies due to characteristics such as age, wealth, and profession have been examined as well.[1]

Non-participation and under-diversification are perhaps the two most examined issues in empirical household investment decisions. Households are more likely to participate if they have higher income, higher financial wealth, higher liabilities, higher education, have disposable income like private pensions or if they are retirees (Calvet, Campbell, and Sodini 2006). Studies seem to suggest greater sophistication in higher

---

[1]Lack of dependence on age in allocation, lower wealth and lower education with greater non-participation and under-diversification, greater sophistication in higher wealth investors (Ameriks and Zeldes (2004)).

wealth investors (Ameriks and Zeldes 2004) where investors with lower wealth and lower education seem to be associated with greater non-participation and underdiversification (Calvet, Campbell, and Sodini 2006). This finding is consistent with previous empirical studies, which demonstrate how wealthy investors invest differently than poorer ones (Tracy, Schneider, and Chan 1999, Heaton and Lucas 2000, and Carroll 2002).

Empirical studies on asset allocation to specific asset classes does provide some limitations. More specifically, Blume and Friend(1975) as well as Kelly (1995) point out that the true picture of asset allocation would require knowledge of mutual fund allocation, to assess an investor's true level of diversification. This point is important because Calvet, Campbell, and Sodini (2006) show that most diversification can be attributed to mutual fund investment, which is consistent with Ameriks and Zeldes' (2003) finding that mutual fund investment dominates stocks in most households. In fact, Calvet, Campbell, and Sodini (2006) find that amongst participating households 87% hold mutual funds where only 55% hold stocks, with 76% of stock holders also owning mutual funds. In terms of asset classes, stocks seem to be dominated by mutual fund strategies, bonds can be seen as safe assets for long term investors (Campbell and Viciera 2002), and a large proportion of investors hold cash. In particular, Calvet, Campbell, and Sodini (2006) document that investors are holding 41% of their assets in cash or cash equivalents (with the remaining wealth distributed around 31% in mutual funds and capital insurance, and 28% in securities). The large exposure to cash may be motivated by the real-life liquidity constraints, that household investors face. It is also important to point out that despite little discussion of real estate investment in traditional portfolio choice, empirical studies show that real estate seems to vary the most drastically based on demographics, of all asset classes, with strong increases in exposure to real estate based on age and wealth (Calvet, Campbell, and Sodini 2006).

Several authors have examined the importance of asset class flows. In particular in a recent paper by Agnew (2003), she asserts that instead of overtrading as was documented in brokerage accounts by Odean (1999) and Odean and Barber (2000),

individual investors actually trade very infrequently. In fact, she finds, by examining asset class flows, that investors often shift out of equities, after extremely negative asset returns, into fixed income products. She claims that this shows that in retirement accounts investors are more prone to exhibit a "flight to safety" instead of explicit return chasing. Given that one in three of the workers in the United States participates in 401(K) programs, it is clear that this "flight to safety" could have a significant impact on market prices as well as demand.

In summary, empirical findings suggest that behavioral effects play an important role in investment decision-making. More specifically, concepts in behavioral finance related to prospect theory, loss aversion, ambiguity aversion, regret theory, lack of self control, and mental accounting provide further insight into why this may be the case.

Chapter 3

Stop-Loss

Stop-loss rules—predetermined policies that reduce a portfolio's exposure after reaching a certain threshold of cumulative losses—are commonly used by retail and institutional investors to manage the risks of their investments, but have also been viewed with some skepticism by critics who question their efficacy. In this paper, we develop a simple framework for measuring the impact of stop-loss rules on the expected return and volatility of an arbitrary portfolio strategy, and derive conditions under which stop-loss rules add or subtract value to that portfolio strategy. We show that under the Random Walk Hypothesis, simple 0/1 stop-loss rules always decrease a strategy's expected return, but in the presence of momentum, stop-loss rules can add value. To illustrate the practical relevance of our framework, we provide an empirical analysis of a stop-loss policy applied to a buy-and-hold strategy in U.S. equities, where the stop-loss asset is U.S. long-term government bonds. Using monthly returns data from January 1950 to December 2004, we find that certain stop-loss rules add 50 to 100 basis points per month to the buy-and-hold portfolio during stop-out periods. By computing performance measures for several price processes, including a new regime-switching model that implies periodic "flights-to-quality", we provide a possible explanation for our empirical results and connections to the behavioral finance literature.

3.1 Introduction

Thanks to the overwhelming dominance of the mean-variance portfolio optimization framework pioneered by Markowitz (1952), Tobin (1958), Sharpe (1964), and Lintner (1965), much of the investments literature—both in academia and in industry—has been focused on constructing well-diversified static portfolios using low-cost index funds. With little use for active trading or frequent rebalancing, this passive perspective comes from the recognition that individual equity returns are difficult to forecast and trading is not costless. The questionable benefits of day-trading are unlikely to outweigh the very real costs of changing one's portfolio weights. It is, therefore, no surprise that a "buy-and-hold" philosophy has permeated the mutual-fund industry and the financial planning profession.[1]

However, this passive approach to investing is often contradicted by human behavior, especially during periods of market turmoil.[2] These behavioral biases sometimes lead investors astray, causing them to shift their portfolio weights in response to significant swings in market indexes, often "selling at the low" and "buying at the high". On the other hand, some of the most seasoned investment professionals routinely make use of systematic rules for exiting and re-entering portfolio strategies based on cumulative losses, gains, and other "technical" indicators.

In this paper, we investigate the efficacy of such behavior in the narrow context of stop-loss rules, i.e., rules for exiting an investment after some threshold of loss is reached and re-entered after some level of gains is achieved. We wish to identify the economic motivation for stop-loss policies so as to distinguish between rational and behavioral explanations for these rules. While certain market conditions may encourage irrational investor behavior—for example, large rapid market declines—

---

[1]This philosophy has changed slightly with the recent innovation of a slowly varying asset allocation that changes according to one's age, e.g., a "lifecycle" fund.

[2]For example, psychologists and behavioral economists have documented the following systematic biases in the human decisionmaking process: overconfidence (Fischoff and Slovic, 1980; Barber and Odean, 2001; Gervais and Odean, 2001), overreaction (DeBondt and Thaler, 1986), loss aversion (Kahneman and Tversky, 1979; Shefrin and Statman, 1985; Odean, 1998), herding (Huberman and Regev, 2001), psychological accounting (Tversky and Kahneman, 1981), miscalibration of probabilities (Lichtenstein et al., 1982), hyperbolic discounting (Laibson, 1997), and regret (Bell, 1982a,b; Clarke et al., 1994).

stop-loss policies are sufficiently ubiquitous that their use cannot always be irrational.

This raises the question we seek to answer in this paper: When do stop-loss rules stop losses? In particular, because a stop-loss rule can be viewed as an overlay strategy for a specific portfolio, we can derive the impact of that rule on the return characteristics of the portfolio. The question of whether or not a stop-loss rule stops losses can then be answered by comparing the expected return of the portfolio with and without the stop-loss rule. If the expected return of the portfolio is higher with the stop-loss rule than without it, we conclude that the stop-loss rule does, indeed, stop losses.

Using simple properties of conditional expectations, we are able to characterize the marginal impact of stop-loss rules on any given portfolio's expected return, which we define as the "stopping premium". We show that the stopping premium is inextricably linked to the stochastic process driving the underlying portfolio's return. If the portfolio follows a random walk, i.e., independently and identically distributed returns, the stopping premium is always negative. This may explain why the academic and industry literature has looked askance at stop-loss policies to date. If returns are unforecastable, stop-loss rules simply force the portfolio out of higher-yielding assets on occasion, thereby lowering the overall expected return without adding any benefits. In such cases, stop-loss rules never stop losses.

However, for non-random-walk portfolios, we find that stop-loss rules *can* stop losses. For example, if portfolio returns are characterized by "momentum" or positive serial correlation, we show that the stopping premium can be positive and is directly proportional to the magnitude of return persistence. Not surprisingly, if conditioning on past cumulative returns changes the conditional distribution of a portfolio's return, it should be possible to find a stop-loss policy that yields a positive stopping premium. We provide specific guidelines for finding such policies under several return specifications: mean reversion, momentum, and Markov regime-switching processes. In each case, we are able to derive explicit conditions for stop-loss rules to stop losses.

Of course, focusing on expected returns does not account for risk in any way. It may be the case that a stop-loss rule increases the expected return but also in-

creases the risk of the underlying portfolio, yielding ambiguous implications for the risk-adjusted return of a portfolio with a stop-loss rule. To address this issue, we compare the variance of the portfolio with and without the stop-loss rule and find that, in cases where the stop-loss rule involves switching to a lower-volatility asset when the stop-loss threshold is reached, the unconditional variance of the portfolio return is reduced by the stop-loss rule. A decrease in the variance coupled with the possibility of a positive stopping premium implies that, within the traditional mean-variance framework, stop-loss rules may play an important role under certain market conditions.

To illustrate the empirical relevance of our analysis, we apply a simple stop-loss rule to the standard asset-allocation problem of stocks vs. bonds using monthly U.S. equity and bond returns from 1950 to 2004. We find that stop-loss rules exhibit significant positive stopping premiums and substantial reductions in variance over large ranges of threshold values—a remarkable feat for a buy-high/sell-low strategy. For example, in one calibration, the stopping premium is approximately 1.0% per annum, with a corresponding reduction in overall volatility of 0.8% per annum, and an average duration of the stopping period of less than 1 month per year. Moreover, we observe conditional-momentum effects following periods of sustained losses in equities that seem to produce scenarios where long-term bonds strongly dominate equities for months at a time. These results suggest that the random walk model is a particularly poor approximation to monthly U.S. equity returns, as Lo and MacKinlay (1999) and others have concluded using other methods.

Motivated by Agnew's (2003) "flight to safety" for household investors, which is similar to the well-documented "flight to quality" phenomenon involving stocks and bonds, we propose a regime-switching model of equity returns in which the Markov regime-switching process is a function of cumulative returns. We show that such a model fits U.S. aggregate stock index data better than other time-series models such as the random walk and AR(1), and can explain a portion of the stopping premium and variance reduction in the historical data.

3.2 Literature Review

Before presenting our framework for examining the performance impact of stop-loss rules, we provide a brief review of the relevant portfolio-choice literature, and illustrate some of its limitations to underscore the need for a different approach.

The standard approach to portfolio choice is to solve an optimization problem in a multi-period setting, for which the solution is contingent on two important assumptions: the choice of objective function and the specification of the underlying stochastic process for asset returns. The problem was first posed by Samuelson (1969) in discrete time and Merton (1969) in continuous time, and solved in both cases by stochastic dynamic programming. As the asset-pricing literature has grown, this paradigm has been extended in a number of important directions.[3]

However, in practice, household investment behavior seems to be at odds with finance theory. In particular, Ameriks and Zeldes (2004) observe that

> ... a great deal of observed variation in portfolio behavior may be explained by the outcome of a few significant decisions that individuals make infrequently, rather than by marginal adjustments continuously.

Moreover, other documented empirical characteristics of investor behavior include non-participation (Calvet, Campbell, and Sodini 2006); under-diversification (Calvet, Campbell, and Sodini 2006); limited monitoring frequency and trading (Ameriks and Zeldes 2004); survival-based selling decisions or a "flight to safety" (Agnew 2003); an absence of hedging strategies (Massa and Simonov, 2004); and concentration in simple strategies through mutual-fund investments (Calvet, Campbell and Sodini 2006). Variations in investment policies due to characteristics such as age, wealth, and profession have been examined as well.[4]

---

[3]For a comprehensive summary of portfolio choice see Brandt (2004). Recent extensions include predictability and autocorrelation in asset returns (Brennan and Xia, 2001; Xia, 2001; Kim and Omberg, 1996; Wachter, 2002; Liu, 1999; and Campbell and Viceria, 1999), model uncertainty (Barberis, 2000), transaction costs (Balduzzi and Lynch, 1999), stochastic opportunity sets (Brennan, Schwartz, and Lagnado, 1997; Brandt, Goyal, Santa-Clara, and Stroud, 2005; and Campbell, Chan, and Viceria, 2003), and behavioral finance (see the references in footnote 2).

[4]For example, lack of age-dependence in allocation, lower wealth and lower education with greater non-participation and under-diversification, and greater sophistication in higher wealth investors

In fact, in contrast to the over-trading phenomenon documented by Odean (1999) and Barber and Odean (2000), Agnew (2003) asserts that individual investors actually trade infrequently. By examining asset-class flows, she finds that investors often shift out of equities after extremely negative asset returns into fixed-income products, and concludes that in retirement accounts, investors are more prone to exhibit a "flight to safety" instead of explicit return chasing. Given that 1 in 3 of the workers in the United States participate in 401(k) programs, it is clear that this "flight to safety" could have a significant impact on market prices as well as demand. Consistent with Agnew's "flight-to-safety" in the empirical application of stop-loss, we find momentum in long-term bonds as a result of sustained periods of loss in equities. This suggests conditional relationships between stocks and bonds, an implication which is also confirmed by our empirical results.[5]

Although stop-loss rules are widely used, the corresponding academic literature is rather limited. The market microstructure literature contains a number of studies about limit orders and optimal order selection algorithms (Easley and O'Hara, 1991; Biais, Hillion, and Spatt, 1995; Chakravarty and Holden, 1995; Handa and Schwartz, 1996; Harris and Hasbrouck, 1996; Seppi, 1997; and Lo, MacKinlay, and Zhang, 2002). Carr and Jarrow (1990) investigate the properties of a particular trading strategy that employs stop-loss orders, and Tschoegl (1988) and Shefrin and Statman (1985) consider behavioral patterns that may explain the popularity of stop-loss rules. However, to date, there has been no systematic analysis of the impact of a stop-loss rule on an existing investment policy, an oversight that we remedy in this paper.

3.3 A Framework for Analyzing Stop-Loss Rules

In this section, we outline a framework for measuring the impact of stop-loss policies on investment performance. In Section 3.3.1, we begin by specifying a simple stop-

---

have all been considered (see Ameriks and Zeldes, 2004).

[5]Although excess performance in long-term bonds may seem puzzling, from a historical perspective, the deregulation of long-term government fixed-income products in the 1950's could provide motivation for the existence of these effects.

loss policy and deriving some basic statistics for its effect on an existing portfolio strategy. We describe several generalizations and qualifications of our framework in Section 3.3.2, and then apply our framework in Section 3.4 to various return-generating processes including the Random Walk Hypothesis, momentum and mean-reversion models, and regime-switching models.

3.3.1 Assumptions and Definitions

Consider any arbitrary portfolio strategy P with returns $\{r_t\}$ that satisfy the following assumptions:

(**A1**) The returns $\{r_t\}$ for the portfolio strategy P are stationary with finite mean μ and variance σ^2.

(**A2**) The expected return μ of P is greater than the riskfree rate r_f, and let $\pi \equiv \mu - r_f$ denote the risk premium of P.

Our use of the term "portfolio strategy" in Assumption (A1) is meant to underscore the possibility that P is a complex dynamic investment policy, not necessarily a static basket of securities. Assumption (A2) simply rules out perverse cases where stop-loss rules add value because the "safe" asset has a higher expected return than the original strategy itself.

Now suppose an investor seeks to impose a stop-loss policy on a portfolio strategy. This typically involves tracking the cumulative return $R_t(J)$ of the portfolio over a window of J periods, where:[6]

$$R_t(J) \equiv \sum_{j=1}^{J} r_{t-j+1} \tag{3.1}$$

and when the cumulative return crosses some lower boundary, reducing the investment in P by switching into cash or some other safer asset. This heuristic approach

---

[6]For simplicity, we ignore compounding effects and define cumulative returns by summing simple returns r_t instead of multiplying $(1+r_t)$. For purposes of defining the trigger of our stop-loss policy, this approximation does not have significant impact. However, we do take compounding into account when simulating the investment returns of a portfolio with and without a stop-loss policy.

motivates the following definition:

Definition 1 *A simple stop-loss policy* $\mathcal{S}(\gamma, \delta, J)$ *for a portfolio strategy* P *with returns* $\{r_t\}$ *is a dynamic binary asset-allocation rule* $\{s_t\}$ *between* P *and a riskfree asset* F *with return* r_f, *where* s_t *is the proportion of assets allocated to* P, *and:*

$$
s_t \;\equiv\; \begin{cases} 0 & \text{if} \quad R_{t-1}(J) \;<\; -\gamma \quad \text{and} \quad s_{t-1} = 1 \quad \text{(exit)} \\ 1 & \text{if} \quad r_{t-1} \;\geq\; \delta \quad \text{and} \quad s_{t-1} = 0 \quad \text{(re-enter)} \\ 1 & \text{if} \quad R_{t-1}(J) \;\geq\; -\gamma \quad \text{and} \quad s_{t-1} = 1 \quad \text{(stay in)} \\ 0 & \text{if} \quad r_{t-1} \;<\; \delta \quad \text{and} \quad s_{t-1} = 0 \quad \text{(stay out)} \end{cases} \tag{3.2}
$$

for $\gamma \geq 0$. *Denote by* r_{st} *the return of portfolio strategy* S, *which is the combinaton of portfolio strategy* P *and the stop-loss policy* $\mathcal{S}$, *hence:*

$$
r_{st} \;\equiv\; s_t r_t \;+\; (1 - s_t) r_f \;. \tag{3.3}
$$

Definition 1 describes a 0/1 asset-allocation rule between P and the riskfree asset F, where 100% of the assets are withdrawn from P and invested in F as soon as the J-period cumulative return $R_{t_1}(J)$ reaches some loss threshold γ at t_1. The stop-loss rule stays in place until some future date $t_2 - 1 > t_1$ when P realizes a return r_{t_2-1} greater than δ, at which point 100% of the assets are transferred from F back to P at date t_2. Therefore, the stop-loss policy $\mathcal{S}(\gamma, \delta, J)$ is a function of three parameters: the loss threshold γ, the re-entry threshold δ, and the cumulative-return window J. Of course, the performance of the stop-loss policy also depends on the characteristics of F—lower riskfree rates imply a more significant drag on performance during periods when the stop-loss policy is in effect.

Note that the specification of the loss and re-entry mechanisms are different; the exit decision is a function of the cumulative return $R_{t-1}(J)$ whereas the re-entry decision involves only the one-period return r_{t-1}. This is intentional, and motivated by two behavioral biases. The first is loss aversion and the disposition effect, in which an individual becomes less risk-averse when facing mounting losses. The second is the "snake-bite" effect, in which an individual is more reluctant to re-enter a portfolio

after experiencing losses from that strategy. The simple stop-loss policy in Definition 1 is meant to address both of these behavioral biases in a systematic fashion.

To gauge the impact of the stop-loss policy $\mathcal{S}$ on performance, we define the following metric:

Definition 2 *The* **stopping premium** $\Delta_\mu(\mathcal{S})$ *of a stop-loss policy $\mathcal{S}$ is the expected return difference between the stop-loss policy $\mathcal{S}$ and the portfolio strategy P:*

$$\Delta_\mu \equiv \mathrm{E}[r_{st}] - \mathrm{E}[r_t] = p_o \left(r_f - \mathrm{E}[r_t|s_t = 0]\right) \qquad (3.4)$$

$$\text{where} \quad p_o \equiv \mathrm{Prob}(s_t = 0) \qquad (3.5)$$

and the **stopping ratio** *is the ratio of the stopping premium to the probability of stopping out:*

$$\frac{\Delta_\mu}{p_o} = r_f - \mathrm{E}[r_t|s_t = 0] . \qquad (3.6)$$

Note that the difference of the expected returns of r_{st} and r_t reduces to the product of the probability of a stop-loss p_o and the conditional expectation of the difference between r_f and r_t, conditioned on being stopped out. The intuition for this expression is straightforward: the only times r_{st} and r_t differ are during periods when the stop-loss policy has been triggered. Therefore, the difference in expected return should be given by the difference in the conditional expectation of the portfolio with and without the stop-loss policy—conditioned on being stopped out—weighted by the probability of being stopped out.

The stopping premium (3.4) measures the expected-return difference per unit time between the stop-loss policy $\mathcal{S}$ and the portfolio strategy P, but this metric may yield misleading comparisons between two stop-loss policies that have very different parameter values. For example, for a given portfolio strategy P, suppose $\mathcal{S}_1$ has a stopping premium of 1% and $\mathcal{S}_2$ has a stopping premium of 2%; this suggests that $\mathcal{S}_2$ is superior to $\mathcal{S}_1$. But suppose the parameters of $\mathcal{S}_2$ implies that $\mathcal{S}_2$ is active only 10% of the time, i.e., 1 month out of every 10 on average, whereas the parameters of

$\mathcal{S}_1$ implies that it is active 25% of the time. On a total-return basis, $\mathcal{S}_1$ is superior, even though it yields a lower expected-return difference per-unit-time. The stopping ratio Δ_μ/p_o given in (3.6) addresses this scale issue directly by dividing the stopping premium by the probability p_o. The reciprocal of p_o is the expected number of periods that $s_t = 0$ or the expected duration of the stop-loss period. Multiplying the per-unit-time expected-return difference Δ_μ by this expected duration $1/p_o$ then yields the total expected-return difference Δ_μ/p_o between r_f and r_t.

The probability p_o of a stop-loss is of interest in its own right because more frequent stop-loss events imply more trading and, consequently, more transactions costs. Although we have not incorporated transactions costs explicitly into our analysis, this can be done easily by imposing a return penalty in (3.3):

$$r_{st} \equiv s_t r_t + (1 - s_t) r_f - \kappa |s_t - s_{t-1}| \tag{3.7}$$

where $\kappa > 0$ is the one-way transactions cost of a stop-loss event. For expositional simplicity, we shall assume $\kappa = 0$ for the remainder of this paper.

Using the metrics proposed in Definition 2, we now have a simple way to answer the question posed in our title: stop-loss policies can be said to stop losses when the corresponding stopping premium is positive. In other words, a stop-loss policy adds value if and only if its implementation leads to an improvement in the overall expected return of a portfolio strategy.

Of course, this simple interpretation of a stop-loss policy's efficacy is based purely on expected return, and ignores risk. Risk matters because it is conceivable that a stop-loss policy with a positive stopping premium generates so much additional risk that the risk-adjusted expected return is less attractive with the policy in place than without it. This may seem unlikely because by construction, a stop-loss policy involves switching out of P into a riskfree asset, implying that P spends more time in higher-risk assets than the combination of P and $\mathcal{S}$. However, it is important to acknowledge that P and $\mathcal{S}$ are dynamic strategies and static measures of risk such as standard deviation are not sufficient statistics for the intertemporal risk/reward trade-offs that

characterize a dynamic rational expectations equilibrium.[7] Nevertheless, it is still useful to gauge the impact of a stop-loss policy on volatility of a portfolio strategy P, as only one of possibly many risk characteristics of the combined strategy. To that end, we have:

Definition 3 *Let the variance difference Δ_{σ^2} of a stopping strategy be given by:*

$$\Delta_{\sigma^2} \equiv \mathrm{Var}[r_{st}] - \mathrm{Var}[r_t] \tag{3.8}$$

$$= \mathrm{E}\big[\mathrm{Var}[r_{st}|s_t]\big] + \mathrm{Var}\big[\mathrm{E}[r_{st}|s_t]\big] - \mathrm{E}\big[\mathrm{Var}[r_t|s_t]\big] - \mathrm{Var}\big[\mathrm{E}[r_t|s_t]\big] \tag{3.9}$$

$$= -p_o \mathrm{Var}[r_t|s_t = 0] +$$
$$p_o(1-p_o)\left[\big(r_f - \mathrm{E}[r_t|s_t = 0]\big)^2 - \left(\frac{\mu - \mathrm{E}[r_t|s_t = 0]}{1 - p_o}\right)^2\right] \tag{3.10}$$

From an empirical perspective, standard deviations are often easier to interpret, hence we also define the quantity $\Delta_\sigma \equiv \sqrt{\mathrm{Var}[r_{st}]} - \sigma$.

Given that a stop-loss policy can affect both the mean and standard deviation of the portfolio strategy P, we can also define the difference between the Sharpe ratios of P with and without $\mathcal{S}$:

$$\Delta_{\mathrm{SR}} \equiv \frac{\mathrm{E}[r_{st}] - r_f}{\sigma_s} - \frac{\mu - r_f}{\sigma}. \tag{3.11}$$

However, given the potentially misleading interpretations of the Sharpe ratio for dynamic strategies such as P and $\mathcal{S}$, we shall refrain from using this metric for evaluating the efficacy of stop-loss policies.[8]

3.3.2 Generalizations and Qualifications

The basic framework outlined in Section 3.3.1 can be generalized in many ways. For example, instead of switching out of P and into a completely riskfree asset, we can allow F to be a lower-risk asset but with some non-negligible volatility. More generally,

---

[7]See Merton (1973) and Lucas (1978), for example.
[8]See Sharpe (1994), Spurgin (2001), and Lo (2002) for details.

instead of focusing on binary asset-allocation policies, we can consider a continuous function $\omega(\cdot) \in [0,1]$ of cumulative returns that declines with losses and rises with gains. Also, instead of a single "safe" asset, we might consider switching into multiple assets when losses are realized, or incorporate the stop-loss policy directly into the portfolio strategy P itself so that the original strategy is affected in some systematic way by cumulative losses and gains. Finally, there is nothing to suggest that stop-loss policies must be applied at the portfolio level—such rules can be implemented security-by-security or asset-class by asset-class.

Of course, with each generalization, the gains in flexibility must be traded off against the corresponding costs of complexity and analytic intractability. These trade-offs can only be decided on a case-by-case basis, and we leave it to the reader to make such trade-offs individually. Our more modest objective in this paper is to provide a complete solution for the leading case of the simple stop-loss policy in Definition (1). From our analysis of this simple case, a number of generalizations should follow naturally, some of which are explored in Kaminski (2006).

However, an important qualification regarding our approach is the fact that we do not derive the simple stop-loss policy (3.2) from any optimization problem—it is only a heuristic, albeit a fairly popular one among many institutional and retail investors. This is a distinct departure from much of the asset-pricing literature in which investment behavior is modelled as the outcome of an optimizing individual seeking to maximize his expected lifetime utility by investing in a finite set of securities subject to a budget constraint, e.g., Merton (1971). While such a formal approach is certainly preferable if the consumption/investment problem is well posed—for example, if preferences are given and the investment opportunity set is completely specified— the simple stop-loss policy can still be studied in the absence of such structure.

Moreover, from a purely behavioral perspective, it is useful to consider the impact of a stop-loss heuristic even if it is not derived from optimizing behavior, precisely because we seek to understand the basis of such behavior. Of course, we can ask the more challenging question of whether the stop-loss heuristic (3.2) can be derived as the optimal portfolio rule for a specific set of preferences, but such inverse-optimal

problems become intractable very quickly (see, for example, Chang, 1988). Instead, we have a narrower set of objectives in this paper: to investigate the basic properties of simple stop-loss heuristics without reference to any optimization problem, and with as few restrictions as possible on the portfolio strategy P to which the stop-loss policy is applied. The benefits of our narrower focus are the explicit analytical results described in Section 3.4, and the intuition that they provide for how stop-loss mechanisms add or subtract value from an existing portfolio strategy.

Although this approach may be more limited in the insights it can provide to the investment process, the siren call of stop-loss rules seems so universal that we hope to derive some useful implications for optimal consumption and portfolio rules from our analysis. Moreover, the idea of overlaying one set of heuristics on top of an existing portfolio strategy has a certain operational appeal that many institutional investors have found so compelling recently, e.g., so-called "portable alpha" strategies. Overlay products can be considered a general class of "superposition strategies", and this is explored in more detail in Kaminski (2006).

3.4 Analytical Results

Having defined the basic framework in Section 3.3 for evaluating the performance of simple stop-loss rules, we now apply them to several specific return-generating processes for $\{r_t\}$, including the Random Walk Hypothesis in Section 3.4.1, mean-reversion and momentum processes in Section 3.4.2, and a statistical regime-switching model in Section 3.4.3. The simplicity of our stop-loss heuristic (3.2) will allow us to derive explicit conditions under which stop-loss policies can stop losses in each of these cases.

3.4.1 The Random Walk Hypothesis

Since the Random Walk Hypothesis is one of the most widely used return-generating processes in the finance literature, any analysis of stop-loss policies must consider this leading case first. Given the framework proposed in Section 3.3, we are able to derive

a surprisingly strong conclusion about the efficacy of stop-loss rules:

Proposition 1 *If $\{r_t\}$ satisfies the Random Walk Hypothesis so that:*

$$r_t = \mu + \epsilon_t \quad , \quad \epsilon_t \overset{\text{IID}}{\sim} \text{White Noise}(0, \sigma_\epsilon^2) \tag{3.12}$$

then the stop-loss policy (3.2) has the following properties:

$$\Delta_\mu = p_o(r_f - \mu) = -p_o\pi \tag{3.13a}$$

$$\frac{\Delta_\mu}{p_o} = -\pi \tag{3.13b}$$

$$\Delta_{\sigma^2} = -p_o\sigma^2 + p_o(1 - p_o)\pi^2 \tag{3.13c}$$

$$\Delta_{\text{SR}} = -\frac{\pi}{\sigma} + \frac{\Delta_\mu + \pi}{\sqrt{\Delta_{\sigma^2} + \sigma^2}} \tag{3.13d}$$

Proof: See Appendix A.2.1. ■

Proposition 1 shows that, for any portfolio strategy with an expected return greater than the riskfree rate r_f, the Random Walk Hypothesis implies that the stop-loss policy (3.2) will always reduce the portfolio's expected return since $\Delta_\mu \leq 0$. In the absence of any predictability in $\{r_t\}$, whether or not the stop-loss is activated has no information content for the portfolio's returns; hence, the only effect of a stop-loss policy is to replace the portfolio strategy P with the riskfree asset when the strategy is stopped out, thereby reducing the expected return by the risk premium of the original portfolio strategy P. If the stop-loss probability p_o is large enough and the risk premium is small enough, (3.13) shows that the stop-loss policy can also reduce the volatility of the portfolio.

The fact that there are no conditions under which the simple stop-loss policy (3.2) can add value to a portfolio with IID returns may explain why stop-loss rules have been given so little attention in the academic finance literature. The fact that the Random Walk Hypothesis was widely accepted in the 1960's and 1970's—and considered to be synonymous with market efficiency and rationality—eliminated the motivation for stop-loss rules altogether. In fact, our simple stop-loss policy may be viewed as a more sophisticated version of the "filter rule" that was tested extensively

by Alexander (1961) and Fama and Blume (1966). Their conclusion that such strategies did not produce any excess profits was typical of the outcomes of many similar studies during this period.

However, despite the lack of interest in stop-loss rules in academic studies, investment professionals have been using such rules for many years, and part of the reason for this dichotomy may be the fact that the theoretical motivation for the Random Walk Hypothesis is stronger than the empirical reality. In particular, Lo and MacKinlay (1988) presented compelling evidence against the Random Walk Hypothesis for weekly U.S. stock-index returns from 1962 to 1985, which has subsequently been confirmed and extended to other markets and countries by a number of other authors. In the next section, we shall see that, if asset-returns do not follow random walks, there are several situations in which stop-loss policies can add significant value to an existing portfolio strategy.

3.4.2 Mean Reversion and Momentum

In the 1980's and 1990's, several authors documented important departures from the Random Walk Hypothesis for U.S. equity returns,[9] and, in such cases, the implications for the stop-loss policy (3.2) can be quite different than in Proposition 1. To see how, consider the simplest case of a non-random-walk return-generating process, the AR(1):

$$r_t = \mu + \rho(r_{t-1} - \mu) + \epsilon_t \ , \quad \epsilon_t \overset{\text{IID}}{\sim} \text{White Noise}(0, \sigma_\epsilon^2) \ , \quad \rho \in (-1, 1) \quad (3.14)$$

where the restriction that ρ lies in the open interval $(-1, 1)$ is to ensure that r_t is a stationary process (see Hamilton, 1994).

This simple process captures a surprisingly broad range of behavior depending on the single parameter ρ, including the Random Walk Hypothesis ($\rho = 0$), mean reversion ($\rho \in (-1, 0)$), and momentum ($\rho = (0, 1)$). However, the implications of this

---

[9]See, for example, Fama and French (1988), Lo and MacKinlay (1988, 1990, 1999), Poterba and Summers (1988), Jegadeesh (1990), Lo (1991), and Jegadeesh and Titman (1993).

return-generating process for our stop-loss rule are not trivial to derive because the conditional distribution of r_t, conditioned on $R_{t-1}(J)$, is quite complex. For example, according to (3.4), the expression for the stopping premium Δ_μ is given by:

$$\Delta_\mu = p_o(r_f - \mathrm{E}[r_t|s_t = 0])$$ (3.15)

but the conditional expectation $\mathrm{E}[r_t|s_t = 0]$ is not easy to evaluate in closed-form for an AR(1). For $\rho \neq 0$, the conditional expectation is likely to differ from the unconditional mean μ since past returns do contain information about the future, but the exact expression is not easily computable. Fortunately, we are able to obtain a good first-order approximation under certain conditions, yielding the following result:

Proposition 2 *If $\{r_t\}$ satisfies an AR(1) (3.14), then the stop-loss policy (3.2) has the following properties:*

$$\frac{\Delta_\mu}{p_o} = -\pi + \rho\sigma + \eta(\gamma, \delta, J)$$ (3.16)

and for $\rho > 0$ and reasonable stop-loss parameters, it can be shown that $\eta(\gamma, \delta, J) \geq 0$, which yields the following lower bound:

$$\frac{\Delta_\mu}{p_o} \geq -\pi + \rho\sigma$$ (3.17)

Proof: See Appendix A.2.2. ∎

Proposition 2 shows that the impact of the stop-loss rule on expected returns is the sum of three terms: the negative of the risk premium, a linear function of the autoregressive parameter ρ, and a remainder term. For a mean-reverting portfolio strategy, $\rho < 0$; hence, the stop-loss policy hurts expected returns to a first-order approximation. This is consistent with the intuition that mean-reversion strategies benefit from reversals, thus a stop-loss policy that switches out of the portfolio after certain cumulative losses will miss the reversal and lower the expected return of the portfolio. On the other hand, for a momentum strategy, $\rho > 0$, in which case there is a possibility that the second term dominates the first, yielding a positive

stopping premium. This is also consistent with the intuition that in the presence of momentum, losses are likely to persist, therefore, switching to the riskfree asset after certain cumulative losses can be more profitable than staying fully invested.

In fact, (3.17) implies that a sufficient condition for a stop-loss policy with reasonable parameters to add value for a momentum-AR(1) return-generating process is

$$\rho \geq \frac{\pi}{\sigma} \equiv \text{SR} \tag{3.18}$$

where SR is the usual Sharpe ratio of the portfolio strategy. In other words, if the return-generating process exhibits enough momentum, then the stop-loss rule will indeed stop losses. This may seem like a rather high hurdle, especially for hedge-fund strategies that have Sharpe ratios in excess of 1.00! However, note that (3.18) assumes that the Sharpe ratio is calibrated at the same sampling frequency as ρ. Therefore, if we are using monthly returns in (3.14), the Sharpe ratio in (3.18) must also be monthly. A portfolio strategy with an annual Sharpe ratio of 1.00—annualized in the standard way by multiplying the monthly Sharpe ratio by $\sqrt{12}$—implies a monthly Sharpe ratio of 0.29, which is still a significant hurdle for ρ but not quite as imposing as 1.00.[10]

3.4.3 Regime-Switching Models

Statistical models of changes in regime, such as the Hamilton (1989) model, are parsimonious ways to capture apparent nonstationarities in the data such as sudden shifts in means and variances. Although such models are, in fact, stationary, they do exhibit time-varying conditional means and variances, conditioned on the particular state that prevails. Moreover, by assuming that transitions from one state to another follow a time-homogenous Markov process, regime-switching models exhibit rich time-series properties that are surprisingly difficult to replicate with traditional linear

---

[10]Of course, the assumption that returns follow an AR(1) makes the usual annualization factor of $\sqrt{12}$ incorrect, which is why we use the phrase "annualized in the standard way". See Lo (2002) for the proper method of annualizing Sharpe ratios in the presence of serial correlation.

processes. Regime-switching models are particularly relevant for stop-loss policies because one of the most common reasons investors put forward for using a stop-loss rule is to deal with a significant change in market conditions such as October 1987 or August 1998. To the extent that this motivation is genuine and appropriate, we should see significant advantages to using stop-loss policies when the portfolio return $\{r_t\}$ follows a regime-switching process.

More formally, let r_t be given by the following stochastic process:

$$r_t = I_t r_{1t} + (1 - I_t) r_{2t} \quad , \quad r_{it} \stackrel{\text{IID}}{\sim} \mathcal{N}(\mu_i, \sigma_i^2) \quad , \quad i = 1, 2 \qquad (3.19\text{a})$$

$$A \equiv \begin{array}{c} \\ I_t = 1 \\ I_t = 0 \end{array} \begin{array}{cc} I_{t+1} = 1 & I_{t+1} = 0 \\ \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} \end{array} \qquad (3.19\text{b})$$

where I_t is an indicator function that takes on the value 1 when state 1 prevails and 0 when state 2 prevails, and A is the Markov transition probabilities matrix that governs the transitions between the two states. The parameters of (3.19) are the means and variances of the two states, $(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$, and the transition probabilities (p_{11}, p_{22}). Without any loss in generality, we adopt the convention that state 1 is the higher-mean state so that $\mu_1 > \mu_2$. Given assumption (A2), this convention implies that $\mu_1 > r_f$, which is an inequality we will make use of below. The six parameters of (3.19) may be estimated numerically via maximum likelihood (see, for example, Hamilton, 1994).

Despite the many studies in the economics and finance literatures that have implemented the regime-switching model (3.19), the implications of regime-switching returns for the investment process has only recently been considered (see Ang and Bekaert, 2004). This is due, in part, to the analytical intractability of (3.19)—while the specification may seem simple, it poses significant challenges for even the simplest portfolio optimization process. However, numerical results can easily be obtained via Monte Carlo simulation, and we provide such results in Sections 3.5.

In this section, we investigate the performance of our simple stop-loss policy (3.2)

for this return-generating process. Because of the relatively simple time-series structure of returns within each regime, we are able to characterize the stopping premium explicitly:

Proposition 3 *If $\{r_t\}$ satisfies the two-state Markov regime-switching process (3.19), then the stop-loss policy (3.2) has the following properties:*

$$\Delta_\mu = p_{o,1}(r_f - \mu_1) + p_{o,2}(r_f - \mu_2) \tag{3.20}$$

$$\frac{\Delta_\mu}{p_o} = (1 - \tilde{p}_{o,2})(r_f - \mu_1) + \tilde{p}_{o,2}(r_f - \mu_2) \tag{3.21}$$

where

$$p_{o,1} \equiv \text{Prob}\left(s_t = 0, I_t = 1\right) \tag{3.22a}$$

$$p_{o,2} \equiv \text{Prob}\left(s_t = 0, I_t = 0\right) \tag{3.22b}$$

$$\tilde{p}_{o,2} \equiv \frac{p_{o,2}}{p_o} = \text{Prob}\left(I_t = 0 \mid s_t = 0\right). \tag{3.22c}$$

If the riskfree rate r_f follows the same two-state Markov regime-switching process (3.19), with expected returns r_{f1} and r_{f2} in states 1 and 2, respectively, then the stop-loss policy (3.2) has the following properties:

$$\Delta_\mu = p_{o,1}(r_{f1} - \mu_1) + p_{o,2}(r_{f2} - \mu_2) \tag{3.23}$$

$$\frac{\Delta_\mu}{p_o} = (1 - \tilde{p}_{o,2})(r_{f1} - \mu_1) + \tilde{p}_{o,2}(r_{f2} - \mu_2). \tag{3.24}$$

The conditional probability $\tilde{p}_{o,2}$ can be interpreted as the accuracy of the stop-loss policy in anticipating the low-mean regime. The higher is this probability, the more likely it is that the stop-loss policy triggers during low-mean regimes (regime 2), which should add value to the expected return of the portfolio as long as the riskfree asset-return r_f is sufficiently high relative to the low-mean expected return μ_2.

In particular, we can use our expression for the stopping ratio Δ_μ/p_o to provide a bound on the level of accuracy required to have a non-negative stopping premium. Consider first the case where the riskfree asset r_f is the same across both regimes.

For levels of $\tilde{p}_{o,2}$ satisfying the inequality:

$$\tilde{p}_{o,2} \geq \frac{\mu_1 - r_f}{\mu_1 - \mu_2} \qquad (3.25)$$

the corresponding stopping premium Δ_μ will be non-negative. By convention, $\mu_1 > \mu_2$, and by assumption (A2), $\mu_1 > r_f$, therefore the sign of the right side of (3.25) is positive. If r_f is less than μ_2, then the right side of (3.25) is greater than 1, and no value of $\tilde{p}_{o,2}$ can satisfy (3.25). If the expected return of equities in both regimes dominates the riskfree asset, then the simple stop-loss policy will always decrease the portfolio's expected return, regardless of how accurate it is. To see why, recall that returns are independently and identically distributed within each regime, and we know from Section 3.4.1 that our stop-loss policy never adds value under the Random Walk Hypothesis. Therefore, the only source of potential value-added for the stop-loss policy (3.2) under a regime-switching process is if the equity investment in the low-mean regime has a lower expected return than the riskfree rate, i.e., $\mu_2 < r_f$. In this case, the right side of (3.25) is positive and less than 1, implying that sufficiently accurate stop-loss policies will yield positive stopping premia.

Note that the threshold for positive stopping premia in (3.25) is decreasing in the spread $\mu_1 - \mu_2$. As the difference between expected equity returns in the high-mean and low-mean states widens, less accuracy is needed to ensure that the stop-loss policy adds value. This may be an important psychological justification for the ubiquity of stop-loss rules in practice. If an investor possesses a particularly pessimistic view of the low-mean state—implying a large spread between μ_1 and μ_2—then our simple stop-loss policy may appeal to him even if its accuracy is not very high.

3.5 Empirical Analysis

To illustrate the potential relevance of our framework for analyzing stop-loss rules, we consider the performance of (3.2) when applied to the standard household asset-allocation problem involving just two asset classes: stocks and long-term bonds. Using

monthly stock- and bond-index data from 1950 to 2004, we find that stop-loss policies produce surprising conditional properties in portfolio returns, stopping losses over a wide range of parameter specifications. Our simple stop-loss rule seems to be able to pick out periods in which long-term bonds substantially out-perform equities, which is especially counterintuitive when we consider the unconditional properties of these two asset classes historically.

For our empirical analysis, we use the monthly CRSP value-weighted returns index to proxy for equities and monthly long-term government bond returns from Ibbotson and Associate to proxy for bonds. We also consider Ibbotson's short-term government bond returns for purposes of comparison. Our sample runs from January 1950 to December 2004, the same time span used by Ang and Berkart's (2004) study of regime-switching models and asset allocation. In Section 3.5.4, we consider the longer time span from January 1926 to December 2004 to reduce estimation error for our behavioral regime-switching model estimates.

Asset	Ann. Mean (%)	Ann. SD (%)	ρ_1 (%)	Skew	Kurt	Min (%)	Med (%)	Max (%)	Ann. Sharpe	MDD (%)
Equities	12.5	14.4	2	-0.3	4.7	-21.6	1.3	16.8	0.9	38.4
Long-Term Bonds	6.2	9.0	6	0.6	6.4	-9.8	0.3	15.2	0.7	25.1
Short-Term Bonds	4.8	0.8	96	1.0	4.4	0.0	0.4	1.4	5.8	1.3

Table 3.1: Summary statistics for the CRSP Value-Weighted Total Market Index, and Ibbotson Associates Long-Term and Short-Term Government Bond Indexes, from January 1950 to December 2004.

In Table 3.1, we summarize the basic statistical properties of our dataset. To be consistent with practice, we implement our stop-loss policies using simple returns, but also provide means and standard deviations of log returns for equities and bonds in Table 3.2 to calibrate some of our simulations. The results in Table 3.1 are well known and require little commentary: stocks outperform bonds, long-term bonds outperform short-term bonds, and the corresponding annual volatilities are consistent with the rank-ordering of mean returns.

In Section 3.5.1, we present the performance statistics of our stop-loss policy applied to our stock and bond return series. We provide a more detailed performance attribution of the stop-loss policy in Section 3.5.2. In Section 3.5.3, we compare our empirical findings to simulated results under the Random Walk Hypothesis, momentum and mean reversion, and regime switching. We conclude that stop-loss rules apparently exploit momentum effects in equities and long-term bonds following periods of sustained losses in equities.

3.5.1 Basic Results

The empirical analysis we perform is straightforward: consider investing 100% in equities in January 1950, and apply the simple stop-loss policy (3.2) to this portfolio on a monthly basis, switching to a 100% investment in long-term bonds when stopped out, and switching back into equities 100% when the re-entry threshold is reached. We run this simulation until December 2004, which yields a time series of 660 monthly returns $\{r_{st}\}$ with which we compute the performance statistics in Definition 2.

Specifically, we compute performance measures for the simple stop-loss strategy (3.2) for cumulative-return windows $J = 3$, 6, 12, and 18 months over stop-loss thresholds $\gamma = 4$–14% and re-entry thresholds $\delta = 0\%$ and 2%. The performance measures Δ_μ, Δ_σ, $\frac{\Delta_\mu}{p_o}$, and p_o are graphed in Figure 3-1. Two robust features are immediately apparent: the first is that stopping premiums Δ_μ are positive, and the second is that the volatility differences Δ_σ are also negative. This suggests that stop-loss rules unambiguously add value to mean-variance portfolio optimizers. Moreover, the robustness of the results over a large range of parameter values indicates some significant time-series structure within these two asset classes.

Figure 3-1 also shows that Δ_μ seems to decrease with larger cumulative-return windows, especially for $J = 6$ and 12 months. This finding is consistent with Δ_μ increasing in p_o when the riskfree rate r_f is higher than the conditional expected return of equities, conditioned on being stopped out (see equation (3.15)). For reference, we plot p_o in Figure 3-2.

For reference, we also plot p_o in Figure 3-2 and find that p_o is monotonically
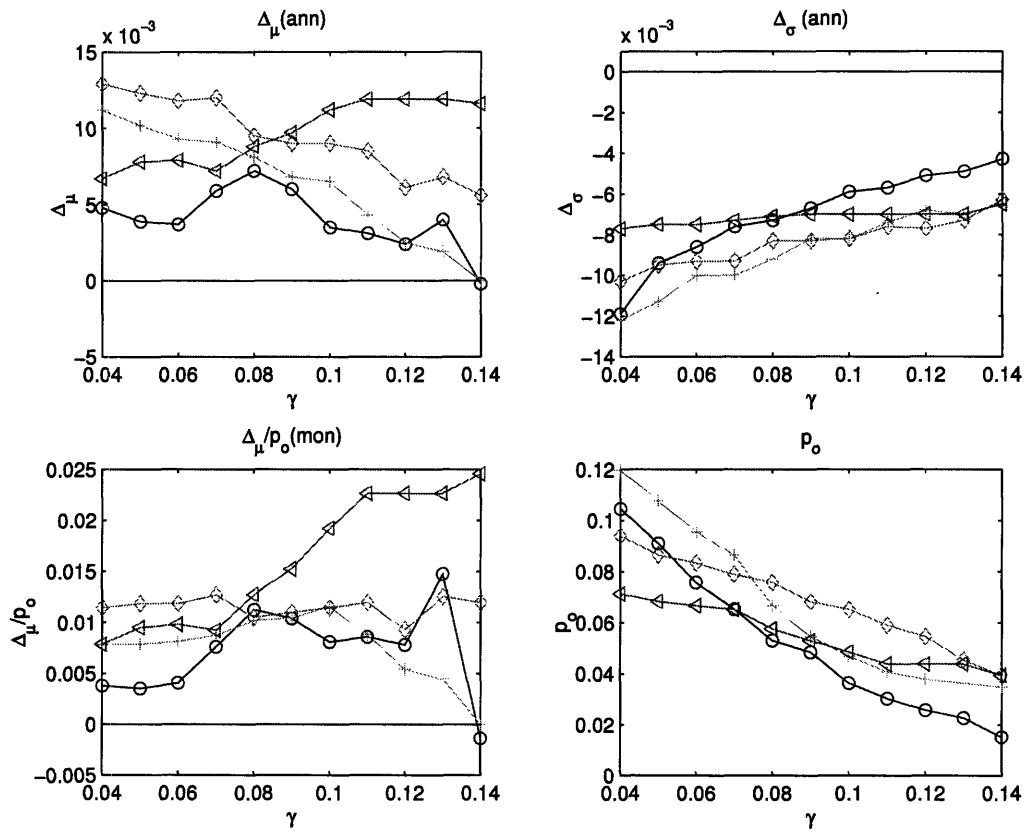
Figure 3-1: Stop-loss performance metrics Δ_μ, Δ_σ, $\frac{\Delta_\mu}{p_o}$, and p_o for the simple stop-loss policy over stopping thresholds $\gamma = 4$–14% with $\delta = 0\%$, $J = 3$ months ($\circ$), 6 months ($+$), 12 months ($\diamond$), and 18 months ($\triangle$).

decreasing with γ as we would expect. In addition, p_o generally ranges between 5% and 10% implying that stop-loss rules stop-out rather infrequently, approximately once a year or once every two years. Nevertheless, these infrequent decisions seem to add considerable value to a buy-and-hold equity portfolio.

Figure 3-1 also contains plots of the stopping ratio Δ_μ/p_o and the figure shows that the stop-loss policy yields an incremental 0.5% to 1% increase in expected return on a monthly basis. The worst Δ_μ/p_o occurs for the 3-month cumulative-return window, and the best Δ_μ/p_o is obtained for large thresholds with an 18-month window size. For the shorter window lengths, smaller thresholds provide less value-added but the value remains positive. However, for the 18-month window, larger thresholds perform better. This connection between stop-loss threshold and cumulative-return window size suggests that there is some fundamental relation—either theoretical or behavioral—between the duration of losses and their magnitude.
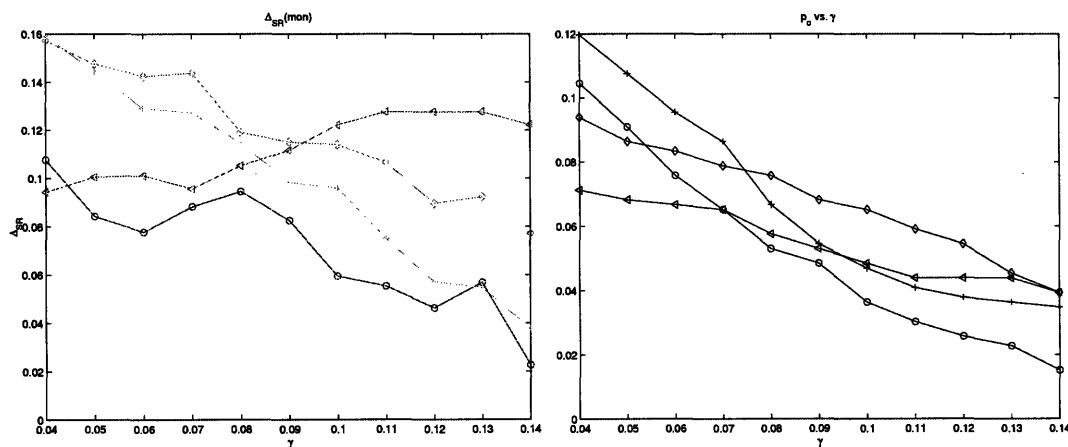
Figure 3-2: Stop-loss performance metrics for Δ_{SR} for the simple stop-loss policy over stopping thresholds $\gamma = 4\text{--}14\%$ with $\delta = 0\%$, $J = 3$ months (o), 6 months (+), 12 months (◇), and 18 months (△).

In Tables A.13 and A.14 of Appendix A.2.4, we examine the performance of equities and bonds during stopped-out periods for stop-loss thresholds $\delta = 0\%$ and $\delta = 2\%$, and find that bonds have significantly better performance with the same level of volatility whereas stocks show reduced performance and increased volatility. We apply a Kolmogorov-Smirnov test to compare the returns before and after stop-loss policies are triggered, and find statistically significant p-values, indicating a difference

between the marginal distribution of returns in and out of stop-out periods (see Table A.15).

Our findings seem to imply momentum-like effects for large negative equity returns, except in the case of large losses over short periods of time which seems to imply reversals. However, since the main focus of our attention is on means and variances, a natural concern is the undue influence of outliers. Even during stop-out periods, we find that the kurtosis of stock and bond returns to be in the range of 2 to 3 (see Tables A.13 and A.14). We also find that the stop-out periods are relatively uniformly distributed over time, refuting the obvious conjecture that a small number of major market crashes are driving the results. Surprisingly, when we exclude the "Tech Bubble" by limiting our sample to December 1999, we find increased performance for our stop-loss policy in most cases. One explanation is that during significant market declines, our stop-loss policy may get back in too quickly, thereby hurting overall performance.

Figure 3-1 also includes a plot of Δ_σ, which shows that volatility is always reduced by the stop-loss policy, but the reduction is smaller for larger stopping thresholds γ. This is to be expected because larger thresholds imply that the stop-loss policy is activated less often. Nevertheless, the reduction in variance is remarkably pronounced for a strategy which so rarely switches out of equities (see Tables A.13 and A.14 for the relative frequency and duration of stop-outs). This reduction seems to be coming from two sources: switching to a lower-volatility asset, and avoiding subsequently higher-volatility periods in equities.

Based on the empirical behavior of Δ_μ and Δ_σ, we expect an increase in the Sharpe ratio, and Figure 3-2 confirms this with a plot of Δ_{SR}. The stop-loss policy has a significant impact on the portfolio's Sharpe ratio even in this simple two-asset case. The relation between Δ_{SR} and window size underscores the potential connection between the amount of time losses are realized and appropriate stop-loss thresholds.

Based on our empirical analysis, we conclude that stop-loss policies could indeed have added value to the typical investor when applied to equities and long-term bonds from 1950 to 2004. In the next two sections, we provide a more detailed analysis of

these results by conducting a performance attribution for the two assets, and by examining several models for asset returns to gauge how substantial these effects are.

3.5.2 Performance Attribution

The empirical success of our simple stop-loss policy implies periods where long-term bonds provide more attractive returns than equities, which beckons us to examine in more detail the properties of both asset classes during stopped-out periods. In particular, we would like to understand if the positive stopping premium is driven by avoiding downside-momentum in equities, positive returns from a flight-to-safety in bonds, or both. Although a closer analysis indicates that both phenomena are present, the conditional performance in bonds seems more compelling. To demonstrate this effect, we examine a specific stop-loss policy and graph the conditional asset-class properties in Figure 3-3, 3-4, and 3-5.

(a)   (b)   (c)
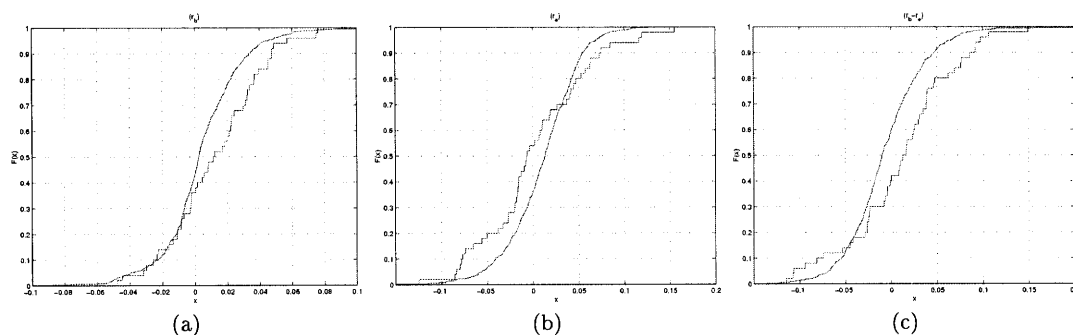
Figure 3-3: Empirical CDFs of (a) Ibbotson Associates Long-Term Government Bond returns (r_b); (b) CRSP Value-Weighted Total Market returns (r_e); (c) and their difference $(r_b - r_e)$, for returns during stopped-out periods (50 data points, dotted line) and non-stopped out periods (610 data points, solid line) with stop-loss parameters $J = 12$, $\gamma = 8\%$, and $\delta = 0\%$, from January 1950 to December 2004.

In Figure 3-3, we plot the empirical cumulative distribution functions (CDFs) for equities, long-term bonds, and their difference for stopped-out and non-stopped-out returns, and in Figure 3-4, we plot the corresponding return histograms for equities and long-term bonds during stopped-out periods, non-stopped-out periods, and both. Figure 3-3 shows that for long-term bonds, returns during stopped-out periods seem to first-order stochastically dominate returns during non-stopped-out period, and that

stopped-out returns exhibit a much larger positive skew. In contrast, equities have larger negative returns and a few larger positive returns, coupled with larger volatility.
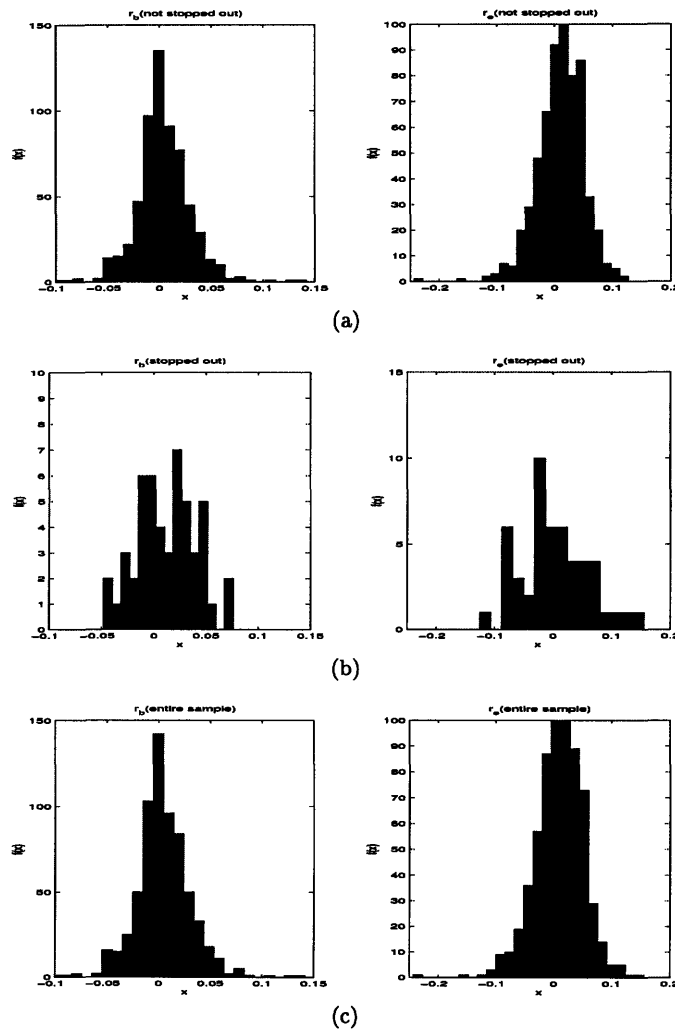
Figure 3-4: Histograms of (a) Ibbotson Associates Long-Term Government Bond returns (r_b); (b) CRSP Value-Weighted Total Market returns (r_e); and (c) their difference $(r_b - r_e)$, for returns during stopped-out periods and the entire sample, with stop-loss parameters $J = 12$, $\gamma = 8\%$, and $\delta = 0\%$, from January 1950 to December 2004.

When we examine the difference between long-term bonds and equities, we find that the CDF of the stopped-out periods almost first-order stochastically dominates the CDF of the non-stopped-out periods, and the positive skew is due to both the increased positive skew in long-term bonds and the large negative returns in equities. The stopped-out difference does not stochastically dominate the non-stopped out periods due to the few large positive returns in equities during stopped-out pe-

riods. By examining these conditional CDFs, we conclude that performance during stopped-out periods is generally good because equities tend to have persistent negative performance and long-term bonds generate excess performance during the periods following negative equity returns. In addition, long-term bonds do not stochastically dominate equities because of the few large reversals in equity returns.
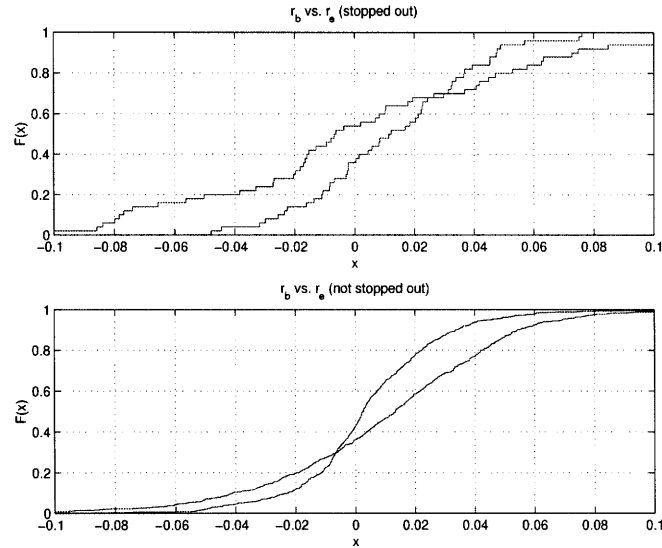
Figure 3-5: Empirical CDFs of Ibbotson Associates Long-Term Government Bond returns (r_b) vs. CRSP Value-Weighted Total Market returns (r_e), for returns during stopped-out periods (50 data points, dotted line) and non-stopped out periods (610 data points, solid line) with stop-loss parameters $J = 12$, $\gamma = 8\%$, and $\delta = 0\%$, from January 1950 to December 2004.

In Figure 3-5, we compare equities to bonds directly by plotting the empirical CDFs for both assets together, for stopped-out and non-stopped-out periods. In this case, we find that during non-stopped-out periods, equities provide a higher return than bonds 70% of the time, but during stopped-out periods, equities provide a higher return only 30% of the time.

3.5.3 A Comparison of Empirical and Analytical Results

To develop further intuition for the empirical results of Section 3.5.1, we conduct several simulation experiments in this section for the return-generating processes of Section 3.4. These simulations will serve as useful benchmarks to gauge the economic

Asset	Return Process	c (%)	k (%)	σ (%)	ρ (%)
Equity	AR(1)	0.93	0.17	4.12	2.52
	AR(1) (ann.)	11.16	2.04	14.28	—
	IID	0.95	0.17	4.12	—
	IID (ann.)	11.46	2.04	14.28	—
Long-Term Bonds	IID	0.48	0.06	2.58	—
	IID (ann.)	5.81	0.80	8.93	—

Table 3.2: Parameter estimates for monthly log returns under both IID and AR(1) return-generating processes for the CRSP Value-Weighted Total Market Index, and IID return-generating process for and Ibbotson Associates Long-Term and Short-Term Government Bond Indexes, from January 1950 to December 2004.

Frequency	μ_{e1} (%)	μ_{e2} (%)	σ_{e1} (%)	σ_{e2} (%)	μ_{b1} (%)	μ_{b2} (%)	σ_{b1} (%)	σ_{b2} (%)	π (%)
Monthly	1.26	0.34	3.11	5.65	0.36	0.72	1.64	3.81	67
Annual	15.14	4.06	10.77	19.57	4.37	8.70	5.67	13.20	—

Table 3.3: Maximum likelihood estimates for a regime-switching model with constant transition probabilities for the CRSP Value-Weighted Total Market return, and Ibbotson Associates Long-Term and Short-Term Government Bond returns, from January 1950 to December 2004.

significance of our empirical results, and can also provide insights into the specific sources of value-added of our stop-loss policy.

We simulate three return-generating processes: the Random Walk Hypothesis, an AR(1) with positive ρ (momentum), and the regime-switching model (3.19). For each process, we simulate 10,000 histories of artificial equity and bond return series, each series containing 660 normally distributed monthly returns (the same sample size as our data), and calibrated to match the means and standard deviations of our data. The parameter estimates used for the IID and AR(1) cases are given in Table 3.2, and the regime-switching parameter estimates, estimated by maximum likelihood, are given in Table 3.3.

For each return history, we apply our stop-loss policy (3.2), compute the perfor-

mance metrics in Definition 2, repeat this procedure 10,000 times, and average the performance metrics across these 10,000 histories. Figure 3-6 plots these simulated metrics for the three return-generating processes, along with the empirical performance metrics for the stop-loss policy with a window size $J = 12$ months and a re-entry threshold of 0%.

Given our analysis of the Random Walk Hypothesis in Section 3.4.1, it is clear that IID returns will yield a negative stopping premium. According to Proposition 1, we know the value of the stopping premia Δ_μ depends on our choice of stopping threshold only through p_o, and the value of $\frac{\Delta_\mu}{p_o} = r_f - \mu$ is constant. Figure 3-6 confirms these implications, and also illustrates the gap between the Random Walk simulations and the empirical results which are plotted using the symbol "o". The t-statistics associated with tests that the empirical performance metrics Δ_μ, Δ_σ, and Δ_{SR} are different from their simulated counterparts are all highly significant at the usual levels, implying resounding rejections of the Random Walk Hypothesis. Alternatively, for our simulations to be consistent with our empirical findings, long-term bonds would have to earn a premium over equities of approximately 1% per month, and equities would have to have much higher volatility than their historical returns have exhibited.

For the AR(1) simulations, Figure 3-6 shows little improvement in explaining the empirical results with this return-generating process—the simulated stopping premium is still quite negative for the amount of positive autocorrelation we have calibrated according to Table 3.2. Using Proposition 2, we can approximate and bound the value of the stopping ratio to be:

$$\frac{\Delta_\mu}{p_o} \approx r_f - \mu + \rho\sigma = -0.0034$$

which is comparable to the stopping ratio under the Random Walk Hypothesis, -0.0045. Given empirical values for Δ_μ/p_o, we can back out the implied value of ρ under an AR(1); these implied values are given in Table 3.4. Clearly, these implied autocorrelations are unrealistically high for monthly equity returns, suggesting that simple AR(1) momentum cannot explain the empirical success of our stop-loss policy.
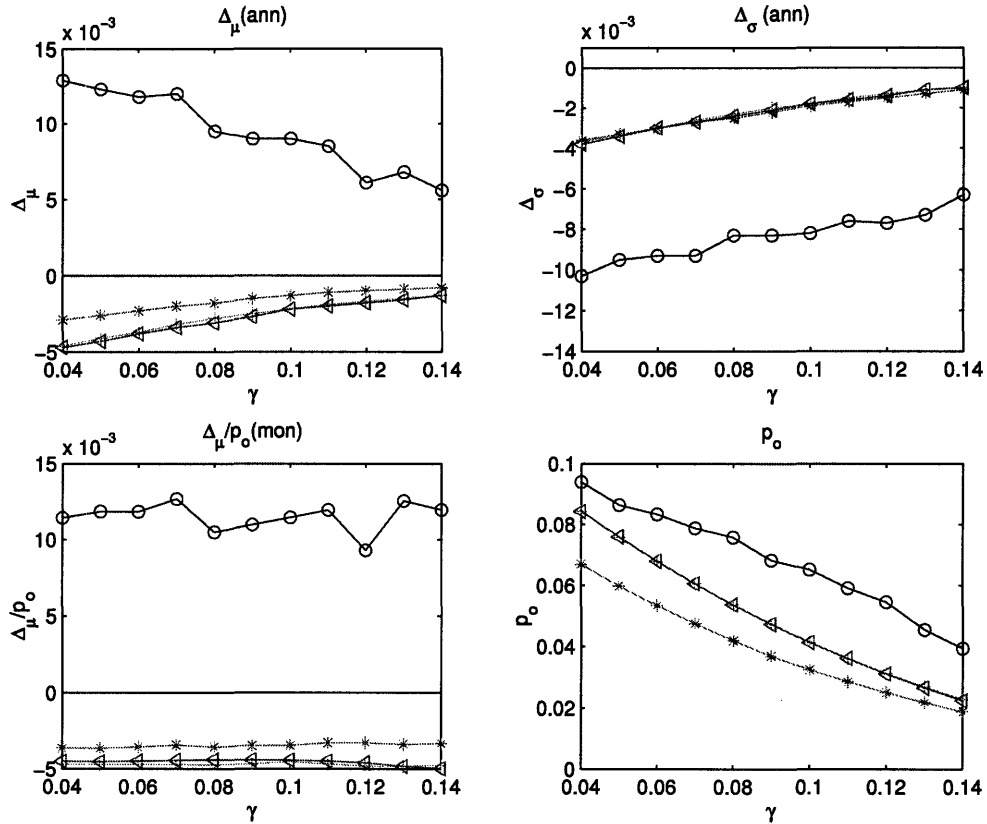
Figure 3-6: Empirical and simulated performance metrics $\Delta_\mu, \Delta_\sigma$, $\frac{\Delta_\mu}{p_o}$, and p_o for the simple stop-loss policy with stopping thresholds $\gamma = 4\text{-}14\%$, $\delta = 0\%$, $J = 12$ months. The empirical results (o) are based on monthly returns of the CRSP Value-Weighted Total Market Index and Ibbotson Associates Long-Term Bond Index from January 1950 to December 2004. The simulated performance metrics are averages across 10,000 replications of 660 monthly normally distributed returns for each of three return-generating processes: IID (+), an AR(1) ($\triangle$), and a regime-switching model ($*$).

J (Months)	Implied ρ (%)	ρ_{MLE} (%)
3	28.1	2.5
6	33.6	2.5
12	39.0	2.5
18	40.1	2.5

Table 3.4: Implied first-order serial correlation coefficient ρ based on the approximation of $\frac{\Delta_\mu}{p_o}$ assuming an AR(1) return-generating process for equities where $\frac{\Delta_\mu}{p_o}$ is an average across the following parameter values for γ: 4%, 5%, 6%, 7%, 8%, 9%, and 10%.

The third set of simulations is based on the regime-switching model (3.19) where long-term bonds are also assumed to vary across regimes, and the parameter estimates in Table 3.3 show some promise of capturing certain features of the data that neither IID nor AR(1) processes can generate. The conditional asymmetry of the two regimes is characterized by one regime with higher returns in equities and lower returns in bonds, and the other with lower returns in equities and higher returns in bonds. Using Proposition 3 (the case with a regime-switching riskfree asset), we can gauge the level of accuracy required of our regime-switching model to obtain a positive stopping premium. Recall from (3.24) that

$$
\begin{aligned}
\frac{\Delta_\mu}{p_o} &= r_{f1} - \mu_1 + \tilde{p}_{o,2}(r_{f2} - r_{f1} + \mu_1 - \mu_2) \\
&= -0.009 + 0.0128\tilde{p}_{o,2}
\end{aligned}
$$

Using this simple result, we see that the stop-loss strategy must correctly switch into bonds with 69.9% accuracy to yield a positive stopping premium. Given the level of volatility in asset returns, it is unrealistic to expect any stopping rule to be able to distinguish regimes with such accuracy. To confirm this intuition, we simulate the regime-switching model using the parameter estimates in Table 3.3 and plot the implied accuracy $\tilde{p}_{o,2}$ over a large range of stop-loss rules in Figure 3-7. The 3-month stopping window outperforms the other stopping windows, especially for large stopping thresholds γ, but none of the implied accuracies comes close to the required accuracy of 69.9% to yield a positive stopping premium. Despite the intuitive appeal of the regime-switching model, it cannot easily account for the empirical success of our simple stop-loss policy.

3.5.4   A Behavioral Regime-Switching Model

Given the lack of success in the regime-switching model (3.19) to explain the empirical performance of the simple stop-loss policy, we propose an alternative based on the flight-to-safety phenomenon. The motivation for such an alternative is the mounting empirical and experimental evidence that investors have two modes of behavior: a
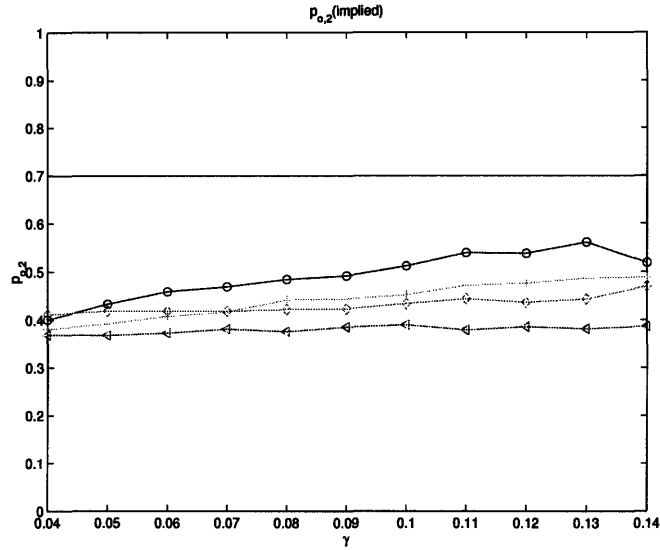
Figure 3-7: The probability of correctly being out of equities during the low-mean regime for equities, or $\tilde{p}_{o,2}$, under Hamilton's (1989) regime-switching model for Basic Stop Loss Rules over various stopping thresholds $\gamma = 4\text{--}14\%$ with $\delta = 0\%$, $J = 3$ (o), 6 (+), 9 (◇), and 18 (△) months, threshold for positive Δ_μ.

normal state, and a distressed or panic state.[11] An implication of this behavior is that investors are asymmetrically impacted by losses, resulting in a flight to safety. The "distress state" is characterized by a lower mean in equities, as well as a higher mean in bonds, and one possible trigger is a sufficiently large cumulative decline in an investor's wealth, e.g., a 401(k) account (Agnew, 2003)

This phenomenon can be captured parsimoniously by extending the regime-switching model (3.19) to allow the regime-switching probabilities to be time-varying and dependent on a cumulative sum of past asset returns:

$$\text{Prob}\,(\,I_t{=}0\,|\,I_{t-1}{=}1\,) \;=\; \frac{\exp(a_1 + b_1 R_{t-1}(J))}{1 + \exp(a_1 + b_1 R_{t-1}(J))} \tag{3.26a}$$

$$\text{Prob}\,(\,I_t{=}1\,|\,I_{t-1}{=}0\,) \;=\; \frac{\exp(a_1 + b_1 R_{t-1}(J))}{1 + \exp(a_2 + b_2 R_{t-1}(J))}\;. \tag{3.26b}$$

The motivation for such a specification is to capture the flight-to-safety effect where the probability of switching to the distress state increases as cumulative losses mount,

---

[11]Examples of such evidence include: disposition effects (Shefrin and Statman, 1985; Odean, 1998, 1999); disappointment aversion (Gul, 1991); loss aversion and prospect theory (Kahneman and Tversky, 1979,1992); and regret (Bell, 1982a,b; Loomes and Sugden, 1982).

which implies a negative b_1 coefficient if we continue to adopt the convention that state 1 is the higher-mean state.[12] This behavioral regime-switching model can be estimated via maximum likelihood estimation following an approach similar to Ang and Bekaert (2004) (see Appendix A.2.3 for details), and the parameter estimates for our monthly equity and long-term bond return series are given in Table 3.5. With the exception of the case where $J = 18$, the b_1 coefficient estimates are indeed negative, consistent with the flight-to-safety phenomenon. Moreover, the coefficient estimates b_2 are positive and much larger in absolute value than the b_1 estimates, implying a stronger tendency to return to the high-mean state from the low-mean state given a cumulative gain of the same absolute magnitude. The fact that both b_1 and b_2 estimates are the largest in absolute value for the shortest horizon $J = 3$ is also consistent with the behavioral evidence that losses and gains concentrated in time have more salience than those over longer time periods.

J	μ_{e1} (%)	μ_{e2} (%)	σ_{e1} (%)	σ_{e2} (%)	μ_{b1} (%)	μ_{b2} (%)	σ_{b1} (%)	σ_{b2} (%)	a_1	a_2	b_1	b_2	σ_{eb1}	σ_{eb2}
Monthly:														
3	1.05	0.32	3.43	5.82	0.33	0.82	1.90	3.87	-4.02	-5.00	-7.53	24.05	0.24	0.15
6	1.04	0.40	3.42	5.68	0.35	0.73	1.85	3.82	-3.87	-4.04	-3.00	10.10	0.22	0.16
12	1.03	0.36	3.41	5.69	0.34	0.76	1.85	3.83	-3.52	-3.14	-2.99	2.47	0.23	0.16
18	1.08	0.48	3.27	5.46	0.34	0.79	1.73	3.64	-4.51	-3.95	4.25	5.47	0.00	0.01
Annual:														
3	12.59	3.78	11.89	20.17	3.90	9.84	6.56	13.41	-4.02	-5.00	-7.53	24.05	0.24	0.15
6	12.48	4.76	11.85	19.67	4.17	8.74	6.41	13.25	-3.87	-4.04	-3.00	10.10	0.22	0.16
12	12.31	4.33	11.81	19.71	4.07	9.09	6.41	13.26	-3.52	-3.14	-2.99	2.47	0.23	0.16
18	12.94	5.73	11.32	18.90	4.04	9.48	5.99	12.59	-4.51	-3.95	4.25	5.47	0.00	0.01

Table 3.5: Maximum likelihood estimates of the behavioral regime-switching model for monthly and annual log-returns for the CRSP Value-Weighted Total Market Index and Ibbotson Associates Long-Term Government Bond Index, from January 1950 to December 2004, and for cumulative-return windows $J = 3$, 6, 12, and 18 months.

Using the maximum likelihood estimates in Table 3.5, we can compute the implied accuracy $\tilde{p}_{o,2}$ required to achieve a positive stopping premium, and these thresholds are given in Table 3.6. These more plausible thresholds—for example, 58.9% for

---

[12]According to (3.26a), a negative value for b_1 implies that cumulative losses would increase the probability of transitioning from state 1 to state 2.

3-month returns—show that a regime-switching model, modified to include time-varying transition probabilities based on cumulative returns, is capable of explaining the empirical results of Section 3.5. Moreover, a simulation experiment similar to those of Section 3.5.3, summarized in Table 3.7, also yields levels of implied accuracy levels required to yield positive stopping premia.

J (Months)	Bound on $p_{o,2} \Rightarrow \Delta_\mu \geq 0$
3	58.9
6	67.6
12	63.4
18	70.4

Table 3.6: Implied lower bound for the accuracy $\tilde{p}_{o,2}$ of the simple stop-loss policy to ensure a positive stopping premia, based on maximum likelihood estimates of the behavioral regime-switching model applied to monthly returns of the CRSP Value-Weighted Total Market Index and Ibbotson Associates Long-Term Government Bond Index, from January 1950 to December 2004.

These results confirm the intuition that regime-switching models—properly extended to incorporate certain behavioral features—can explain more of the empirical performance of simple stop-loss rules than the other return-generating processes we have explored. In fact, the differences between the empirical and simulated performance of our stop-loss policy are not statistically significant under the behavioral regime-switching model for many of the stop-loss parameters, and the behavioral regime-switching model generates variance patterns that are more consistent with those in the data.

However, despite providing a better explanation of the empirical success of our stop-loss policy, the behavioral regime-switching model cannot generate the magnitude of stopping premia observed in the historical record. Therefore, stop-loss policies must be exploiting additional time-varying momentum in equities and long-term bonds that we have not completely captured in our time-series models of stock and bond returns. We leave this as a direction for future research.

J (Months)	n	Implied Empirical $p_{o,2}$	Implied Range of $p_{o,2}$ with Simulation	Bound on $p_{o,2}$ $\Rightarrow \Delta_\mu \geq 0$
3	3	1.16	[0.421,0.834]	0.589
	6	1.58	[0.404,0.714]	0.676
	12	1.70	[0.341,0.566]	0.634
	18	1.84	[0.432,0.583]	0.704
6	3	1.16	[0.485,0.847]	0.589
	6	1.58	[0.499,0.762]	0.676
	12	1.70	[0.418,0.604]	0.634
	18	1.84	[0.474,0.619]	0.704
12	3	1.16	[0.578,0.785]	0.589
	6	1.58	[0.565,0.713]	0.676
	12	1.70	[0.486,0.668]	0.634
	18	1.84	[0.528,0.593]	0.704
18	3	1.16	[0.663,0.857]	0.589
	6	1.58	[0.635,0.871]	0.676
	12	1.70	[0.545,0.604]	0.634
	18	1.84	[0.594,0.691]	0.704

Table 3.7: Simulated values for implied $\tilde{p}_{o,2}$, and thresholds for positive stopping premium based on maximum likelihood parameter estimates of the behavioral regime-switching model with behavioral cumulative-return windows of length n and stop-loss cumulative-return windows of length J.

3.6 Conclusion

In this paper, we provide an answer to the question when do stop-loss rules stop losses? The answer depends, of course, on the return-generating process of the underlying investment for which the stop-loss policy is implemented, as well as the particular dynamics of the stop-loss policy itself. If "stopping losses" is interpreted as having a higher expected return with the stop-loss policy than without it, then for a specific binary stop-loss policy, we derive various conditions under which the expected-return difference—which we call the stopping premium—is positive. We show that under the most common return-generating process–the Random Walk Hypothesis—the stopping premium is always negative. The widespread cultural affinity for the Random Walk Hypothesis, despite empirical evidence to the contrary, may explain the general indifference to stop-loss policies in the academic finance literature.

However, under more empirically plausible return-generating processes such as momentum or regime-switching models, we show that stop-loss policies can gener-

ate positive stopping premia. And when applied to the standard household asset-allocation decision between U.S. equities and long-term bonds from January 1950 to December 2004, we find a substantially positive stopping premium with a correspondingly large reduction in variance. These empirical results suggest important nonlinearities in aggregate stock and bond returns that have not been fully explored in the empirical finance literature. For example, our analysis suggests elevated levels of momentum associated with large negative returns, and asymmetries in asset returns following periods of cumulative losses.

Our analytical and empirical results contain several points of intersection with the behavioral finance literature. First, the flight-to-safety phenomena—best illustrated by events surrounding the default of Russian government debt in August 1998—may create momentum in equity returns and increase demand for long-term bonds, creating positive stopping premia as a result. Second, systematic stop-loss policies may profit from the disposition effect and loss aversion, the tendency to sell winners too soon and hold on to losers too long. Third, if investors are ambiguity-averse, large negative returns may cause them to view equities as more ambiguous which, in relative terms, will make long-term bonds seem less ambiguous. This may cause investors to switch to bonds to avoid uncertainty about asset returns.

More generally, there is now substantial evidence from the cognitive sciences literature that losses and gains are processed by different components of the brain. These different components provide a partial explanation for some of the asymmetries observed in experimental and actual markets. In particular, in the event of a significant drop in aggregate stock prices, investors who are generally passive will become motivated to trade because mounting losses will cause them to pay attention when they ordinarily would not. This influx of uninformed traders, who have less market experience and are more likely to make irrational trading decisions, can have a significant impact on equilibrium prices and their dynamics. Therefore, even if markets are usually efficient, on occasions where a significant number of investors experience losses simultaneously, markets may be dominated temporarily by irrational forces. The mechanism for this coordinated irrationality is cumulative loss.

Of course, our findings shed little light on the controversy between market efficiency and behavioral finance. The success of our simple stop-loss policy may be due to certain nonlinear aspects of stock and bond returns from which our strategy happens to benefit, e.g., avoiding momentum on the downside and exploiting asymmetries in asset returns following periods of negative cumulative returns. And from the behavioral perspective, our stop-loss policy is just one mechanism for avoiding or anticipating the usual pitfalls of human judgment, e.g., the disposition effect, loss aversion, ambiguity aversion, and flight-to-safety.

In summary, both behavioral finance and rational asset-pricing models may be used to motivate the efficacy of stop-loss policies, in addition to the widespread use of such policies in practice. This underscores the importance of learning how to deal with loss as an investor, of which a stop-loss rule is only one dimension. As difficult as it may be to accept, for the millions of investors who lamented after the bursting of the Technology Bubble in 2000 that "if I only got out earlier, I wouldn't have lost so much", they may have been correct.

Chapter 4

General Superposition Strategies

Superposition type stopping rules - or predetermined policies to stop out of underlying portfolio strategies - are commonly used by investors to get in and out of positions. The most well-known type of this superposition or overlay type strategy is that of the stop-loss rule from the previous chapter. Given the analysis of stop-loss, in this chapter I examine a more general approach to overlay type strategies and discuss how the performance of superposition strategies can be evaluated in general terms. This general framework is based on the classic investor's utility maximization problem and suggests how stopping rules can impact arbitrary portfolio strategies. By relaxing the assumption that the true stochastic process for asset returns is known, stopping rules can exploit higher order properties of asset returns by solving alternate optimization problems. This approach can be specialized to the simple case of buy-and-hold mean variance preferences to allow for a discussion of the random walk hypothesis versus predictability.

4.1 Background

When you talk to investors, of any type, whether he or she is an investment practitioner, hedge fund manager, or a household investor, it remains clear that all investors are fundamentally concerned with the act of getting in and out of investments. This observation is clearly supported by the principles of behavioral finance laid out by

prospect theory, loss aversion, regret aversion, etc. As a result, when investors make decisions, they are fundamentally concerned about relative performance and the realizing of losses and gains. They also use mental accounting principles which cause them to disaggregate and aggregate their performance. As a result, instead of performing an aggregate optimization, investors use rules and heuristics consistent with human decision making.

Recent extensions to the portfolio choice problem have highlighted the importance of perturbations to the investor's problem.[1] These extensions demonstrate how portfolio strategies can still remain optimal while departing from a traditional myopic portfolio strategy as proposed by Merton (1969, 1971) and Samuelson (1969). In contrast with the theory proposed by portfolio choice, the empirical evidence, as well as, a review of the financial planning industry, demonstrates that investors adopt simple heuristics for investment. Consistent with common anecdotal evidence from investors and the principles of mental accounting, these simple heuristics are often supplemented by simple rules and policies which help investors to engage in the act of buying and selling. Although the practice of adopting simple rules to allow for self control may seem inconsistent with models which call for optimization, the practice of using heuristics and rules is hardwired into the human decision-making process. This observation highlights the fact that the optimization problem abstracts away from real difficulty investors face with the physical act of buying and selling. Based on the principles of human decision-making, it is clear that humans adopt and fine tune simple decision rules to allow which make discrete decisions.

Anecdotal evidence from investors and empirical evidence in portfolio performance supports the conclusion that investors use simple rules and heuristics. Ameriks and Zeldes (2004) demonstrate that while using simple heuristic strategies, it is the discrete decisions investor's make that creates the largest determinant of portfolio performance. It is precisely these discrete decisions which are of interest in this chapter.

---

[1]Predictability (Brennen and Xia (2001), Xia (2001)), non-standard utility function (Ait-Sahalia and Brandt (2001)), portfolio constraints (Balduzzi and Lynch (1999)), and model uncertainty (Barberis (2000))

4.2 Framework for Analyzing Superposition Stopping Rules

In this section, I outline a framework for measuring the impact of superposition type stopping policies on investment performance. In Section 4.2.1, I begin by specifying a class of stopping policies and quantify their impact on the simple Investor's problem. I outline the specific case of maximizing expected return with two assets under buy-and-hold strategies in Section 4.3. I then apply my framework to discuss two classic stopping rule examples including the stop-loss and buy-low sell-high strategies in Section 4.4.

4.2.1 Assumptions and Definitions

Consider an arbitrary portfolio strategies P_a with $a \in A$, where A is a set of arbitrary portfolio strategies and stopping policies $\mathcal{S} \in S$, where S is a set of stopping policies. Any arbitrary portfolio strategy P_a has returns $\{r_{at}\}$ and satisfies the following assumptions:

(**A1**) The returns $\{r_{at}\}$ for the portfolio strategy P_a are stationary with finite mean $\mu_a(f)$ and variance $\sigma_a^2(f)$ under any stationary distribution for asset returns f.

(**A2**) For all arbitrary portfolio strategies $a \in A$, where A is the set of arbitrary portfolio strategies - the addition of any stopping rule $\mathcal{S} \in S$ does not alter the arbitrary portfolio strategy - the stopping policy is superimposed onto the arbitrary portfolio strategy P_a.

(**A3**) There exists a true stochastic process f^* for asset returns $\{r_t\}$ which is stationary.

My use of the term "portfolio strategy" in Assumption (A1) is meant to underscore the possibility that P is a complex dynamic investment policy, not necessarily a static basket of securities. Assumption (A2) maintains that a stopping policy is superimposed onto an arbitrary portfolio strategy with out altering the dynamics

of the arbitrary strategy. This assumption is crucial for defining an "overlay" or "superposition" type strategy and in the definition of the performance metrics for the stopping policy. Assumption (A3) is important for benchmarking the performance of arbitrary portfolio strategies. The key point of this analysis is that f^* exists but is unknown. My use of the expression $\{r_t\}$ is meant to underscore the possibility that $\{r_t\}$ is a vector of asset returns not necessarily one risky asset. From this point on, I can assume for notational purposes that given a utility function $U(r_{at})$, defined over the asset returns $\{r_{at}\}$ of a portfolio strategy, is maximized under the true distribution (f^*) of asset returns by arbitrary portfolio strategy P_{a^*}. Consistent with classic mean variance analysis, I formulate the one period investor's problem for a given utility function. The investor maximizes his or her utility over portfolio strategies under the true distribution for asset returns f^*.

Definition 4 *The Simple Investor's Problem is to find the best portfolio strategy $a \in A$ which maximizes the expected utility $\mathrm{E}[U(r_{at})]$ as a function of the portfolio strategy P_a with returns $\{r_{at}\}$ where the true distribution for asset returns $\{r_t\}$ is f^*.*

$$\max_{a \in A} \mathrm{E}[U(r_{at})] \tag{4.1}$$

$$\{r_t\} \sim f^* \tag{4.2}$$

Remaining consistent with industry practice and mental accounting principles, I then define a stopping policy. Under Assumption (A2), the stopping policy does not alter the dynamics of an underlying strategy. The key difference between an arbitrary portfolio strategy and a stopping policy is that a stopping policy is applied to a portfolio strategy; it is only a rule. On the other hand, the combination of an underlying strategy and a stopping rule creates a different portfolio strategy P_s. Motivated by the use of stopping rules in practice, it is precisely this new strategy which is of interest in this section.

Definition 5 *A stopping policy $\mathcal{S}(\Gamma, F_{t-1})$ for a underlying portfolio strategy P_u with returns $\{r_{ut}\}$ is a dynamic binary asset-allocation rule $\{s_t\}$ between P_u and an*

alternative arbitrary portfolio strategy P_a with returns $\{r_{at}\}$, where s_t is the proportion of assets allocated to P_u, and:

$$s_t \equiv f(\mathcal{S}(\Gamma, F_{t-1})) \tag{4.3}$$

s_t is determined by the stopping policy defined by $\mathcal{S}(\Gamma, F_{t-1})$ where F_{t-1} is the information vector at time $t-1$ and Γ is a vector of parameters. Denote by $\{r_{st}\}$ the return of overall new portfolio strategy P_s, which is the combination of underlying portfolio strategy P_u and the alterative portfolio strategy P_a with $a \in A$ using the corresponding stopping policy $\mathcal{S} \in S$, hence:

$$r_{st} \equiv s_t r_{ut} + (1 - s_t) r_{at} . \tag{4.4}$$

Letting P_u be the investor's choice of underlying strategy with returns $\{r_{ut}\}$, I can simply assume that P_u is the strategy which is the optimal portfolio strategy for the Investor's problem assuming asset returns follow an approximation of the stochastic process f_u. Thus,

$$\mathrm{E}_{f_u}[U(r_{ut})] \geq \mathrm{E}_{f_u}[U(r_{at})] \quad \forall a \in A \tag{4.5}$$

For an arbitrary portfolio strategy P_u, the investor can improve his or her expected utility by maximizing over a set of stopping rules $\mathcal{S} \in S$ and alternative strategies P_a forall $a \in A$. This results in a new portfolio strategy P_s. This option for improving the portfolio strategy suggests a new performance metric $\Delta_U(s)$, which is outlined in the following definition.

Definition 6 *The **stopping utility premium** $\Delta_U(s)$ of portfolio strategy P_s with returns $\{r_{st}\}$ for an underlying portfolio strategy P_u with stopping policy $\mathcal{S} \in S$ into an alternative portfolio strategy P_a with returns $\{r_{at}\}$ with $a \in A$ is the expected difference in utility between the underlying portfolio strategy P_u and the new portfolio*

strategy P_s with the stopping policy:

$$\Delta_U(s) \equiv \mathrm{E}[U(r_{st})] - \mathrm{E}[U(r_{ut})] = p_o \left(\mathrm{E}[U(r_{at}) - U(r_{ut})|s_t = 0] \right) \quad (4.6)$$

$$\text{where} \quad p_o \equiv \mathrm{Prob}(s_t = 0) \quad (4.7)$$

Given this definition the expected utility of the new portfolio strategy P_s is the following:

$$\mathrm{E}[U(r_{st})] = \mathrm{E}[U(r_{ut})] + \Delta_U(s) \quad (4.8)$$

Using the definition of the utility premium, I can define a new optimization problem for the investor.

Definition 7 *Given an underlying portfolio strategy P_u with returns $\{r_{ut}\}$, the Investor's Stopping Policy Problem is to find a new portfolio strategy P_s using a stopping policy $\mathcal{S} \in S$ and alternative portfolio strategy P_a with returns $\{r_{at}\}$ to maximizes the stopping utility premium $\Delta_U(s)$.*

$$\max_{\mathcal{S} \in S, a \in A} \Delta_U(s) \quad (4.9)$$

$$\{r_t\} \sim f^* \quad (4.10)$$

4.2.2 Policy Improvement and Uncertainty

Given that the true stochastic process for asset returns f^* is unknown, I can then clarify how a stopping rule can improve the performance of an underlying strategy. If I first consider the case that the stochastic process is known, the option to use stopping policies can clearly improve the performance of an arbitrary portfolio strategy $a \in A$.

Corollary 8 *Given an underlying portfolio strategy P_u with returns $\{r_{ut}\}$, if there exists an alternative portfolio strategy P_a with returns $\{r_{at}\}$ and stopping policy $\mathcal{S} \in S$ such that a and $\mathcal{S}$ are a solution to the Investor's stopping policy Problem with $\Delta_U(s) > 0$ for new portfolio strategy P_s then $\mathrm{E}_{f^*}[U(r_{st})] > \mathrm{E}_{f^*}[U(r_{ut})]$ and P_s dominates underlying portfolio strategy P_u.*

Proof: *Using the definition of* $\Delta_U(s)$,

$$
\begin{aligned}
\mathrm{E}_{f^*}[U(r_{st})] &= \mathrm{E}_{f^*}[U(r_{ut})] + \Delta_U(s) \\
&> \mathrm{E}_{f^*}[U(r_{ut})]
\end{aligned}
$$

■

This result is trivial for the case where (f^*) is known, because in this case the Investor will choose a^*. For any other arbitrary portfolio strategy the investor will choose the stopping policy $s_t = 1$ with an alternative portfolio strategy P_{a^*} and stop out for all time into P_{a^*}. Despite the simplicity of the case where the distribution is known, this result still gives intuition into how a stopping policy may improve an arbitrary portfolio strategy, when I relax my assumptions about the distribution of asset returns. As in reality, an investor must assume approximations to the true stochastic process and design his or her portfolio strategy P based on these assumptions. Given both the complicated nature of parameter estimation, model dependence, and the intractability in finding portfolio strategies under complicated stochastic processes, investors generally assume rather simple stochastic processes which best fit the aggregate properties of asset returns. As a result, I return to the underlying portfolio strategy P_u. An investor, who assumes that f_u is a good aggregate approximation to f^*, will select portfolio strategy P_u, where P_u is the solution to the Investor's problem under f_u. In addition to the investor's beliefs about aggregate return dynamics, an investor may also examine other approximations to the true stochastic process denoted by $f_{u'}$. The motivation for the use of $f_{u'}$ is that $f_{u'}$ may be a better conditional approximation of f^*, but not necessarily a better aggregate approximation of f^*. In practice, as models of asset returns become increasingly complex to fit higher order or state dependent dynamics of asset returns, there is a substantial tradeoff in overspecification, model uncertainty, and parameter estimation problems. These problems can outweigh the benefits of finding optimal policies (See Barberis 2000). As a result, the investor only uses the approximation $f_{u'}$ to make conditional portfolio decisions. I outline how the use of conditional approximations to f^* can be connected

to the use of stopping policies in the following corollaries. In Corollary 9, I relax the assumption that f^* is known but allow the conditional approximation to be exact over a subset of policies and alternative strategies. In Corollary 10, I extend this by allowing the conditional approximation of f^* to be exact within some bound, over a subset of policies and alternative strategies.

Corollary 9 *Given an underlying portfolio strategy P_u and there exists a subset of stopping policies $\mathcal{S} \in \bar{S} \subset S$ and arbitrary portfolio strategies $a \in \bar{A} \subset A$ such that conditioned on the stopping policy being activated (i.e. $s_t = 0$), $\Delta_U(s)|_{f_{u'}} = \Delta_U(s)|_{f^*}$ for the new portfolio strategy P_s. If there exists a solution $(\mathcal{S}, a)$ for portfolio strategy P_s to the Investor's stopping problem with $\Delta_U(s) > 0$ under $f_{u'}$ then*

$$\mathrm{E}_{f^*}[U(r_{st})] > \mathrm{E}_{f^*}[U(r_{ut})] \tag{4.11}$$

Proof:

$$
\begin{aligned}
\mathrm{E}_{f^*}[U(r_{st})] &= \mathrm{E}_{f^*}[U(r_{ut})] + \Delta_U(s')|_{f^*} & (4.12)\\
&= \mathrm{E}_{f^*}[U(r_{ut})] + \Delta_U(s')|_{f_{u'}} & (4.13)\\
&> \mathrm{E}_{f^*}[U(r_{ut})] & (4.14)
\end{aligned}
$$

∎

This corollary explains how an alternative stochastic process $f_{u'}$, which fits the conditional properties of the true stochastic process f^*, can suggest stopping policies that improve the underlying choice of portfolio strategy. A key difference between this corollary and traditional literature in portfolio choice, is that it does not depend on one aggregate stochastic process. In reality, it may be the case that some stochastic processes are conditionally a better fit.

Corollary 10 *Given an underlying portfolio strategy P_u and there exists a subset of stopping policies $\mathcal{S} \in \bar{S} \subset S$ and arbitrary portfolio strategies $a \in \bar{A} \subset A$ such that under an approximate stochastic process $f_{u'}$ the approximate stochastic process*

is ϵ approximate in utility function over all $\mathcal{S} \in \bar{S}$ and $a \in \bar{A}$ for resulting portfolio strategies P_s, hence:

$$|\Delta_U(s)|_{f^*} - \Delta_U(s)|_{f_{u'}}| \leq \epsilon \quad \forall \mathcal{S} \in \bar{S}, a \in \bar{A}$$

If there exists a solution to the Investor's stopping policy Problem $(\mathcal{S}, a) \in (\bar{S}, \bar{A})$ for portfolio strategy P_s with $\Delta_U(s)|_{f_{u'}} > \epsilon$ then P_s dominates P_u under f^*.

$$\mathrm{E}_{f^*}[U(r_{st})] > \mathrm{E}_{f^*}[U(r_{ut})] \tag{4.15}$$

and hence;

$$|\mathrm{E}_{f^*}[U(r_{a^*t})] - \mathrm{E}_{f^*}[U(r_{st})]| < |\mathrm{E}_{f^*}[U(r_{a^*t})] - \mathrm{E}_{f^*}[U(r_{ut})]| \tag{4.16}$$

Proof: *Given that* $|\Delta_U(s)|_{f^*} - \Delta_U(s)|_{f_{u'}}| \leq \epsilon$ *if* $\Delta_U(s)|_{f_{u'}} > \epsilon$ *then* $\Delta_U(s)|_{f_{u'}} > 0$ *and*

$$\mathrm{E}_{f^*}[U(r_{st})] = \mathrm{E}_{f^*}[U(r_{ut})] + \Delta_U(s')|_{f^*} \tag{4.17}$$

$$> \mathrm{E}_{f^*}[U(r_{ut})] \tag{4.18}$$

The previous corollary explains how under $f_{u'}$, a stopping policy can be used to improve overall performance of a strategy.

4.2.3 Utility Premiums

Consistent with my analysis of stop-loss rules in Chapter 3, I discuss specific cases of utility functions, define general notions of stopping premiums, and define other performance metrics. These definitions include the stopping premium, variance difference, volatility difference, and Sharpe Ratio difference.

Definition 11 *Given an underlying portfolio strategy P_u with returns $\{r_{ut}\}$, the stopping premium $\Delta_\mu(s)$ of a new portfolio strategy P_s with stopping policy $\mathcal{S} \in S$ and alternative portfolio strategy P_a with returns $\{r_{at}\}$ is the expected return difference*

between the underlying portfolio strategy P_u and the new portfolio strategy P_s:

$$\Delta_\mu(s) \equiv \mathrm{E}[r_{st}] - \mathrm{E}[r_{ut}] = p_o\left(\mathrm{E}[r_{at} - r_{ut}|s_t = 0]\right) \qquad (4.19)$$

$$\text{where} \quad p_o \equiv \mathrm{Prob}(s_t = 0) \qquad (4.20)$$

and the **stopping ratio** *is the ratio of the stopping premium to the probability of stopping out:*

$$\frac{\Delta_\mu}{p_o}(s) = \mathrm{E}[r_{at} - r_{ut}|s_t = 0] . \qquad (4.21)$$

Note that the difference of the expected returns of r_{st} and r_{ut} reduces to the product of the probability of a stopping out of the underlying strategy p_o and the conditional expectation of the difference between alternative portfolio strategy r_{at} and the underlying portfolio strategy r_{ut}, conditioned on being stopped out. The intuition for this expression is straightforward: the only times r_{st} and r_{ut} differ are during periods when the stopping policy has been triggered. Therefore, the difference in expected return should be given by the difference in the conditional expectation of the portfolio with and without the stopping policy—conditioned on being stopped out—weighted by the probability of being stopped out.

The stopping premium (4.19) measures the expected-return difference per unit time between the new stopping strategy P_s with stopping policy $\mathcal{S}$ and alternative portfolio strategy P_a and the underlying portfolio strategy P_u, but this metric may yield misleading comparisons between two stopping policies that have very different parameter values. For example, for a given underlying portfolio strategy P_u, suppose $\mathcal{S}_1$ has a stopping premium of 1% and $\mathcal{S}_2$ has a stopping premium of 2%; this suggests that $\mathcal{S}_2$ is superior to $\mathcal{S}_1$. But suppose the parameters of $\mathcal{S}_2$ implies that $\mathcal{S}_2$ is active only 10% of the time, i.e., 1 month out of every 10 on average, whereas the parameters of $\mathcal{S}_1$ implies that it is active 25% of the time. On a total-return basis, $\mathcal{S}_1$ is superior, even though it yields a lower expected-return difference per-unit-time. The stopping ratio Δ_μ/p_o given in (4.21) addresses this scale issue directly by dividing the stopping

premium by the probability p_o. The reciprocal of p_o is the expected number of periods that $s_t = 0$ or the expected duration of the stopping period. Multiplying the per-unit-time expected-return difference Δ_μ by this expected duration $1/p_o$ then yields the total expected-return difference Δ_μ/p_o between r_{at} and r_{ut}.

Of course, this simple interpretation of a stopping policy's efficacy is based purely on expected return, and ignores risk. Risk matters because it is conceivable that a stopping policy with a positive stopping premium generates so much additional risk that the risk-adjusted expected return is less attractive with the policy in place than without it. However, it is important to acknowledge that P_u, P_a, P_s and $\mathcal{S}$ are dynamic strategies and static measures of risk such as standard deviation are not sufficient statistics for the intertemporal risk/reward trade-offs that characterize a dynamic rational expectations equilibrium.[2] Nevertheless, it is still useful to gauge the impact of a stopping policy on volatility of an underlying portfolio strategy P_u, as only one of possibly many risk characteristics of the combined strategy. To that end, I have:

Definition 12 *Let the variance difference $\Delta_{\sigma^2}(s)$ of a stopping strategy $\mathcal{S}$ for new portfolio strategy P_s be given by:*

$$\Delta_{\sigma^2}(s) \equiv \operatorname{Var}[r_{st}] - \operatorname{Var}[r_{ut}] \tag{4.22}$$

$$= \operatorname{E}\big[\operatorname{Var}[r_{st}|s_t]\big] + \operatorname{Var}\big[\operatorname{E}[r_{st}|s_t]\big] \tag{4.23}$$

$$- \operatorname{E}\big[\operatorname{Var}[r_{ut}|s_t]\big] - \operatorname{Var}\big[\operatorname{E}[r_{ut}|s_t]\big] \tag{4.24}$$

From an empirical perspective, standard deviations are often easier to interpret, hence I also define the quantity $\Delta_\sigma(s) \equiv \sqrt{\operatorname{Var}[r_{st}]} - \sigma$.

Given that a stopping policy can affect both the mean and standard deviation of the underlying portfolio strategy P_u, I can also define the difference between the Sharpe ratios of P_u with and without $\mathcal{S}$:

Definition 13 *Let the Sharpe ratio difference $\Delta_{\mathrm{SR}}(s)$ of a stopping strategy P_s be*

---

[2]See Merton (1973) and Lucas (1978), for example.

given by:

$$\Delta_{\text{SR}}(s) \equiv \frac{\text{E}[r_{st}] - r_f}{\sigma_s} - \frac{\text{E}[r_{ut}] - r_f}{\sigma} . \qquad (4.25)$$

In addition to basic utility premiums, stochastic dominance may also provide a method for determining if a stopping policy will add value to an underlying strategy. Stochastic dominance theory and empirical tests for stochastic dominance are summarized in Section A.3.

4.3　A Simple Case: Two Assets, Buy-and-Hold

Having defined the basic framework in Section 4.2 for evaluating the performance of general "superposition" type stopping rules, I now discuss the implications for the simple case of maximizing expected return over one risky and one riskfree asset. This analysis parallels my analysis of stop-loss in Chapter 3. I first outline how this scenario simplifies the performance metrics defined in Section 4.2.3. I examine several data return generating processes including the Random Walk Hypothesis in Section 4.3.1, general predictability in Section 4.3.2, and a statistical regime-switching model in Section 4.3.3. Using the simple definition for the stopping premium $\Delta_\mu(s)$, I can then outline explicit conditions under which stopping rules can add value to the underlying buy-and-hold strategy in the risky asset.

Proposition 4 *Let the underlying portfolio strategy P_u be a buy-and-hold strategy in the risky asset $\{r_t\}$, then a stopping policy S into an alternative portfolio strategy P_a where P_a is a buy-and-hold strategy in the riskfree asset $\{r_f\}$. The following*

simplifications for the performance metrics hold:

$$\Delta_\mu(s) \equiv p_o \left(r_f - \mathrm{E}[r_t|s_t = 0] \right) \tag{4.26}$$

$$\frac{\Delta_\mu(s}{p_o} \equiv r_f - \mathrm{E}[r_t|s_t = 0] \tag{4.27}$$

$$\Delta_{\sigma^2}(s) \equiv -p_o \mathrm{Var}[r_t|s_t = 0] \; +$$
$$p_o(1 - p_o) \left[\left(r_f - \mathrm{E}[r_t|s_t = 0] \right)^2 \; - \; \left(\frac{\mu - \mathrm{E}[r_t|s_t = 0]}{1 - p_o} \right)^2 \right] \tag{4.28}$$

$$\Delta_\sigma(s) \equiv \sqrt{\sigma^2 + \Delta_{\sigma^2}(s)} - \sigma \tag{4.29}$$

$$\text{where} \quad p_o \equiv \mathrm{Prob}(s_t = 0) \tag{4.30}$$

4.3.1 The Random Walk Hypothesis

Since the Random Walk Hypothesis is one of the most widely used return-generating processes in the finance literature, any analysis of stopping policies must consider this leading case first. Given the framework proposed in Section 4.2, I am able to derive a surprisingly strong conclusion about the efficacy of stopping rules which is consistent with the findings of Alexander (1961) and Fama and Blume (1966):

Proposition 5 *If $\{r_t\}$ satisfies the Random Walk Hypothesis so that:*

$$r_t = \mu + \epsilon_t \;, \quad \epsilon_t \overset{\mathrm{IID}}{\sim} \text{White Noise}(0, \sigma_\epsilon^2) \tag{4.31}$$

Given two buy-and-hold strategies then any stopping policy S for the stopping strategy P_s has the following properties:

$$\Delta_\mu(s) = p_o(r_f - \mu) = -p_o \pi \tag{4.32a}$$

$$\frac{\Delta_\mu}{p_o}(s) = -\pi \tag{4.32b}$$

$$\Delta_{\sigma^2}(s) = -p_o \sigma^2 + p_o(1 - p_o)\pi^2 \tag{4.32c}$$

$$\Delta_{\mathrm{SR}}(s) = -\frac{\pi}{\sigma} + \frac{\Delta_\mu + \pi}{\sqrt{\Delta_{\sigma^2} + \sigma^2}} \tag{4.32d}$$

Proposition 5 shows that, for any portfolio strategy with an expected return greater than the riskfree rate r_f, the Random Walk Hypothesis implies that any stopping

policy will always reduce the portfolio's expected return since $\Delta_\mu(s) \leq 0 \; \forall s$. In the absence of any predictability in $\{r_t\}$, whether or not a stopping policy is activated, has no information content for the portfolio's returns; hence, the only effect of a stopping policy is to replace the underlying portfolio strategy P_u in the risky asset with the riskfree asset when the strategy is stopped out, thereby reducing the expected return by the risk premium of the original portfolio strategy P_u. If the stopped out probability p_o is large enough and the risk premium is small enough, (4.32) shows that the stopping policy can also reduce the volatility of the portfolio.

The fact that there are no conditions under which any stopping policy can add value to a buy-and-hold portfolio with IID returns may explain why stopping rules such as stop-loss have been given so little attention in the academic finance literature. The fact that the Random Walk Hypothesis was widely accepted in the 1960's and 1970's—and considered to be synonymous with market efficiency and rationality— eliminated the motivation for stopping rules altogether. In fact, these stopping policies may be viewed as a more sophisticated version of the "filter rule" that was tested extensively by Alexander (1961) and Fama and Blume (1966). Their conclusion that such strategies did not produce any excess profits was typical of the outcomes of many similar studies during this period.

However, despite the lack of interest in stopping rules in academic studies, investment professionals have been using such rules for many years, and part of the reason for this dichotomy may be the fact that the theoretical motivation for the Random Walk Hypothesis is stronger than the empirical reality. In particular, Lo and MacKinlay (1988) presented compelling evidence against the Random Walk Hypothesis for weekly U.S. stock-index returns from 1962 to 1985, which has subsequently been confirmed and extended to other markets and countries by a number of other authors. In the next section, I demonstrate that, if asset-returns do not follow random walks, there are many situations in which stopping policies can add significant value to an existing underlying portfolio strategy.

4.3.2 General Predictability

In the 1980's and 1990's, several authors documented important departures from the Random Walk Hypothesis for U.S. equity returns,[3] and, in such cases, the implications for the stopping policies can be quite different than in Proposition 5. In this simple case, if I examine the stopping ratio $\frac{\Delta_\mu}{p_o}(s)$ I see that I have a condition for when the stopping premium will become positive.

Proposition 6 *If there exists a stopping policy S which satisfies the following, then* $\Delta_\mu(s) \geq 0$

$$E[\pi_t|s_t = 0] \leq 0 \quad \Rightarrow \quad \Delta_\mu(s) \geq 0 \qquad (4.33)$$

Proposition (6) highlights how a stopping rule can exploit conditional predictability by finding times when the conditional risk premium is negative when conditioned on the stopping rule. This is demonstrated more clearly in the following section on regime switching processes and a clearer example is provided using stop-loss in Chapter 3.

4.3.3 Regime-Switching Models

Statistical models of changes in regime, such as the Hamilton (1989) model, are parsimonious ways to capture apparent nonstationarities in data such as sudden shifts in means and variances. Although such models are, in fact, stationary, they do exhibit time-varying conditional means and variances, conditioned on the particular state that prevails. Moreover, by assuming that transitions from one state to another follow a time-homogenous Markov process, regime-switching models exhibit rich time-series properties that are surprisingly difficult to replicate with traditional linear processes. To the extent that this motivation is genuine and appropriate, I examine the efficacy of stopping rules for the particular case when the portfolio return $\{r_t\}$ follows a regime-switching process.

---

[3]See, for example, Fama and French (1988), Lo and MacKinlay (1988, 1990, 1999), Poterba and Summers (1988), Jegadeesh (1990), Lo (1991), and Jegadeesh and Titman (1993).

More formally, let r_t be given by the following stochastic process:

$$r_t = I_t r_{1t} + (1 - I_t) r_{2t} \quad , \quad r_{it} \overset{\text{IID}}{\sim} \mathcal{N}(\mu_i, \sigma_i^2) \quad , \quad i = 1, 2 \qquad (4.34a)$$

$$A \equiv \begin{matrix} & I_{t+1}=1 & I_{t+1}=0 \\ I_t=1 & \\ I_t=0 \end{matrix} \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} \qquad (4.34b)$$

where I_t is an indicator function that takes on the value 1 when state 1 prevails and 0 when state 2 prevails, and A is the Markov transition probabilities matrix that governs the transitions between the two states. The parameters of (4.34) are the means and variances of the two states, $(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$, and the transition probabilities (p_{11}, p_{22}). Without any loss in generality, I adopt the convention that state 1 is the higher-mean state so that $\mu_1 > \mu_2$. If I assume the aggregate risk premium is non-negative, this implies that $\mu_1 > r_f$, which is an inequality I will make use of below. The six parameters of (4.34) may be estimated numerically via maximum likelihood (see, for example, Hamilton, 1994).

Proposition 7 *If $\{r_t\}$ satisfies the two-state Markov regime-switching process (4.34), then a stopping policy $\mathcal{S}$ has the following properties:*

$$\Delta_\mu(s) = p_{o,1}(r_f - \mu_1) + p_{o,2}(r_f - \mu_2) \qquad (4.35)$$

$$\frac{\Delta_\mu}{p_o}(s) = (1 - \tilde{p}_{o,2})(r_f - \mu_1) + \tilde{p}_{o,2}(r_f - \mu_2) \qquad (4.36)$$

where

$$p_{o,1} \equiv \text{Prob}(s_t = 0, I_t = 1) \qquad (4.37a)$$

$$p_{o,2} \equiv \text{Prob}(s_t = 0, I_t = 0) \qquad (4.37b)$$

$$\tilde{p}_{o,2} \equiv \frac{p_{o,2}}{p_o} = \text{Prob}(I_t = 0 \mid s_t = 0) . \qquad (4.37c)$$

If the riskfree rate r_f follows the same two-state Markov regime-switching process (4.34), with expected returns r_{f1} and r_{f2} in states 1 and 2, respectively, then the

stopping policy (S) has the following properties:

$$\Delta_\mu(s) = p_{o,1}(r_{f1} - \mu_1) + p_{o,2}(r_{f2} - \mu_2) \tag{4.38}$$

$$\frac{\Delta_\mu}{p_o}(s) = (1 - \tilde{p}_{o,2})(r_{f1} - \mu_1) + \tilde{p}_{o,2}(r_{f2} - \mu_2) . \tag{4.39}$$

The conditional probability $\tilde{p}_{o,2}$ can be interpreted as the accuracy of a stopping policy in anticipating the low-mean regime. The higher is this probability, the more likely it is that the stop-loss policy triggers during low-mean regimes (regime 2), which should add value to the expected return of the portfolio as long as the riskfree asset-return r_f is sufficiently high relative to the low-mean expected return μ_2.

In particular, I can use the expression for the stopping ratio $\Delta_\mu/p_o(s)$ to provide a bound on the level of accuracy required to have a non-negative stopping premium. Consider first the case where the riskfree asset r_f is the same across both regimes. For levels of $\tilde{p}_{o,2}$ satisfying the inequality:

$$\tilde{p}_{o,2} \geq \frac{\mu_1 - r_f}{\mu_1 - \mu_2} \tag{4.40}$$

the corresponding stopping premium $\Delta_\mu(s)$ will be non-negative. By convention, $\mu_1 > \mu_2$, and if I assume that the aggregate risk premium is non-negative, $\mu_1 > r_f$, therefore the sign of the right side of (4.40) is positive. If r_f is less than μ_2, then the right side of (4.40) is greater than 1, and no value of $\tilde{p}_{o,2}$ can satisfy (4.40). If the expected return of equities in both regimes dominates the riskfree asset, then any stopping policy will always decrease the portfolio's expected return, regardless of how accurate it is. To see why, recall that returns are independently and identically distributed within each regime, and we know from Section 4.3.1 that the stopping policy never adds value under the Random Walk Hypothesis. Therefore, the only source of potential value-added for a stopping policy (S) under a regime-switching process is if the equity investment in the low-mean regime has a lower expected return than the riskfree rate, i.e., $\mu_2 < r_f$. In this case, the right side of (4.40) is positive and less than 1, implying that sufficiently accurate stopping policies will yield positive

stopping premia.

Note that the threshold for positive stopping premia in (4.40) is decreasing in the spread $\mu_1 - \mu_2$. As the difference between expected equity returns in the high-mean and low-mean states widens, less accuracy is needed to ensure that a stopping policy adds value. This may be an important psychological justification for the ubiquity of so called stopping rules like stop-loss in practice. If an investor possesses a particularly pessimistic view of the low-mean state—implying a large spread between μ_1 and μ_2— then any such stopping policy may appeal to him even if its accuracy is not very high.

The conclusions in this section exactly mirror those discussed in my study of stop-loss. Most importantly, all stopping rules can be connected to the accuracy of the stopping rule based on the asymmetries in asset regimes. More specifically, the accuracy required for using a stopping rule depends explicitly on the level of asymmetry in asset returns in vary regimes. Thus, consistent with practice, if an investor believes that a rule has sufficient accuracy in predicting regime shifts, he or she may consider applying this rule. Classic examples include simple trend following rules based on moving averages, buy-high sell-low strategies, stop-loss rules, stop-gain rules, etc. As a result, this link between stopping rules and accuracy confirm the intuition behind the classic saying, that a good investor knows "when to get in and when to get out."

4.4 Examples of Stopping Policies

In this section, I discuss a few popular examples of stopping rules including stop-loss and buy-low sell-high. I demonstrate how these stopping rules fit into the definitions I proposed in Section 4.2.1. The stop-loss rule was examined in detail in Chapter 3, and I discuss the results for the buy-low sell-high strategy in Section 4.4.2.

4.4.1 Stop-Loss

Investor's commonly apply stop-loss rules to take losses in investments. This typically involves tracking state dependent quantities such as a cumulative return $R_t(J)$ of the portfolio over a window of J periods, where:[4]

$$R_t(J) \equiv \sum_{j=1}^{J} r_{t-j+1} \tag{4.41}$$

and when the cumulative return crosses some lower boundary, reducing the investment in P by switching into cash or some other safer asset. This heuristic approach motivates the following definition:

Definition 14 *A* **simple stop-loss policy** $\mathcal{S}(\gamma, \delta, J)$ *for a portfolio strategy P with returns $\{r_t\}$ is a dynamic binary asset-allocation rule $\{s_t\}$ between P and a riskfree asset F with return r_f, where s_t is the proportion of assets allocated to P, and:*

$$s_t \equiv \begin{cases} 0 & \text{if} & R_{t-1}(J) & < & -\gamma & \text{and} & s_{t-1} = 1 & \text{(exit)} \\ 1 & \text{if} & r_{t-1} & \geq & \delta & \text{and} & s_{t-1} = 0 & \text{(re-enter)} \\ 1 & \text{if} & R_{t-1}(J) & \geq & -\gamma & \text{and} & s_{t-1} = 1 & \text{(stay in)} \\ 0 & \text{if} & r_{t-1} & < & \delta & \text{and} & s_{t-1} = 0 & \text{(stay out)} \end{cases} \tag{4.42}$$

for $\gamma \geq 0$.

Definition 14 describes a 0/1 asset-allocation rule between P and the riskfree asset F, where 100% of the assets are withdrawn from P and invested in F as soon as the J-period cumulative return $R_{t_1}(J)$ reaches some loss threshold γ at t_1. The stop-loss rule stays in place until some future date $t_2 - 1 > t_1$ when P realizes a return r_{t_2-1} greater than δ, at which point 100% of the assets are transferred from F back to P at date t_2. Therefore, the stop-loss policy $\mathcal{S}(\gamma, \delta, J)$ is a function of three parameters: the loss threshold γ, the re-entry threshold δ, and the cumulative-return window J.

---

[4]For simplicity, I ignore compounding effects and define cumulative returns by summing simple returns r_t instead of multiplying $(1+r_t)$. For purposes of defining the trigger of our stop-loss policy, this approximation does not have significant impact. However, I do take compounding into account when simulating the investment returns of a portfolio with and without a stop-loss policy.

The stop-loss policy is examined in further detail in Chapter 3.

4.4.2 Buy-Low Sell-High

Many investment sites enjoy boast about their ability to buy-low and sell-high. Fitting this strategy into the framework of superposition strategies, this typically would involve tracking state dependent quantities such as cumulative returns $R_t(J)$ over a window of J periods. Since the decision to get in and get out may not be symmetric, this strategy may involve tracking cumulative returns over different window sizes as well. When a cumulative window crosses some lower boundary, a buy will be initiated, whereas when a cumulative window crosses some upper boundary a sell with be initiated. This heuristic approach motives the following definition:

Definition 15 *A simple buy-low sell-high policy $\mathcal{S}(\gamma_h, \gamma_l, J_h, J_l)$ for a portfolio strategy P with returns $\{r_t\}$ is a dynamic binary asset-allocation rule $\{s_t\}$ between P and a riskfree asset F with return r_f, where s_t is the proportion of assets allocated to P, and:*

$$
s_t \equiv \begin{cases}
0 & \text{if} \quad R_{t-1}(J_h) > \gamma_h \quad \text{and} \quad s_{t-1} = 1 \quad \text{(sell high)} \\
1 & \text{if} \quad R_{t-1}(J_l) \geq \gamma_l \quad \text{and} \quad s_{t-1} = 0 \quad \text{(buy low)} \\
1 & \text{if} \quad R_{t-1}(J_h) \leq \gamma_h \quad \text{and} \quad s_{t-1} = 1 \quad \text{(stay in)} \\
0 & \text{if} \quad R_{t-1}(J_l) < \gamma_l \quad \text{and} \quad s_{t-1} = 0 \quad \text{(stay out)}
\end{cases}
\tag{4.43}
$$

for $\gamma_h, \gamma_l \geq 0$.

Definition 15 describes a 0/1 asset-allocation rule between P and the riskfree asset F, where 100% of the assets are withdrawn from P and invested in F as soon as the J_h-period cumulative return $R_{t_1}(J)$ reaches some gain threshold γ_h at t_1. The buy low sell high rule stays in place until some future date $t_2 - 1 > t_1$ when P realizes a return $R_{t_2-1}(J_l)$ less than $-\gamma_l$, at which point 100% of the assets are transferred from F back to P at date t_2. Therefore, the buy low sell high policy $\mathcal{S}(\gamma_h, \gamma_l, J_h, J_l)$ is a function of four parameters: the gain threshold γ_h, the loss threshold γ_l, and the cumulative-gain-return window J_h, and the cumulative-loss-return window J_l.

4.5 Discussion and Conclusion

In this chapter, I have examined a class of stopping policies, which I call superposition strategies. These strategies are called superposition strategies because they are simply "overlay" policies, in that they do not alter the underlying dynamics of the strategies they are superimposed on. Since in practice, either based on mental accounting or because of simplicity, investors often choose an underlying strategy. Following this choice, they subsequently perturb this strategy by applying simple rules, which govern when they get in and out of positions. The most common of which is the stop-loss rule, as discussed in the previous chapter. As a result of these perturbations to their underlying strategy, the new strategy they have created what I define as P_s, can be quite different in nature from their original choice. In addition to being consistent with many principles in behavioral finance, this practice is also supported by empirical evidence that suggests that the discrete decisions investors make seem to be the greatest determinant of overall portfolio performance. Given the predominance of stopping rules in practice, the framework I suggest in this chapter, albeit rather simple, lends some direction into understanding how stopping rules can add or subtract value from underlying strategies.

By examining the classic utility framework parallel to mean variance analysis, I define a utility stopping premium. This performance metric allows me to measure the impact of a stopping policy on an underlying strategy. By relaxing assumptions about the true stochastic process for asset returns, I can explain how the use of stopping rules may actually improve the performance of an underlying strategy. This analysis may provide some motivation for saying that successful practitioners seem to "know when to get in and when to get out" of their positions. I show how the appropriate use of stopping rules can actually add value, which is consistent with the theory laid out by portfolio choice. Motivated by the fact that the true return generating process for asset returns is unknown and complex specifications of that return generating process can investors lead astray, the use of simple heuristics and the skillful use of stopping rules can lead to better performance. On the other hand, at the same time,

consistent with classic warnings from financial planning practitioners, even with a simple heuristic strategy, timing the market is extremely difficult, and those investors who advertently or inadvertently apply stopping rules by following market trends risk severely under performing even a simple heuristic strategy.

Chapter 5

Conclusion

Despite the many advances in portfolio choice, a careful look at the financial planning industry and the empirical literature in portfolio choice, demonstrates that despite these advances, financial planning remains more of an "art" than a science. As a result, there is a predominance of simple heuristics and investment rules, both among households, and investment professionals. Despite being at odds with theories of classic utility maximization, the field of behavioral finance lends support to these techniques as being grounded in the basic mechanics of human decision making. As opposed to solving an explicit optimization problem, I focus directly on the application of stopping rules with investment strategies to remain consistent with how investors actually invest. I define a framework for measuring the performance impact of stopping rules consistent with standard portfolio theory.

When I turn my attention to the classic stop-loss rule, I attempt to answer the question "When do stop-loss rules stop losses?", using a simplification of the general framework for superposition type strategies. I show how the answer to this question depends explicitly on the return generating process of the underlying investment, as well as the specific dynamics of the stop-loss policy itself. By defining and examining a stopping premium, I can explicitly determine when stop loss rules may actually add value. In particular, given the most commonly assumed stochastic process, the Random Walk Hypothesis, I show that stop-loss rules never stop losses. In the case of predictability, I show that momentum type effects, modeled using simple serial auto-

correlation and conditional asymmetries in return regimes, can produce a positive stopping premium.

To demonstrate the practical applicability of my approach, I apply my stop-loss rule to the standard household allocation problem between U.S. equities and long-term bonds from January 1950 to December 2004. I find a substantially positive stopping premium, over large ranges of threshold values for the stop-loss rule, which is coupled with substantial reductions of variance. These findings suggest important non-linearities in aggregate stock and bond returns, which may be motivated by some type of flight-to-safety or flight-to-quality.

More generally, there is now substantial evidence from the cognitive sciences literature that losses and gains are processed by different components of the brain. These different components provide a partial explanation for some of the asymmetries observed in experimental and actual markets. In particular, in the event of a significant drop in aggregate stock prices, investors who are generally passive will become motivated to trade, because mounting losses will cause them to pay attention when they ordinarily would not. This influx of uninformed traders, who have less market experience and are more likely to make irrational trading decisions, can have a significant impact on equilibrium prices and their dynamics. Therefore, even if markets are usually efficient, on occasions where a significant number of investors experience losses simultaneously, markets may be dominated temporarily by irrational forces. The mechanism for this coordinated irrationality is cumulative loss.

Of course, these findings shed little light on the controversy between market efficiency and behavioral finance. The success of the simple stop-loss policy may be due to certain nonlinear aspects of stock and bond returns from which the strategy happens to benefit, e.g., avoiding momentum on the downside and exploiting asymmetries in asset returns following periods of negative cumulative returns. And from the behavioral perspective, the stop-loss policy is just one mechanism for avoiding, or anticipating, the usual pitfalls of human judgment, e.g., the disposition effect, loss aversion, ambiguity aversion, and flight-to-safety.

In summary, both behavioral finance and rational asset-pricing models may be

used to motivate the efficacy of stop-loss policies, supporting the widespread use of such policies in practice. This underscores the importance of learning how to deal with loss as an investor, of which a stop-loss rule is only one dimension. As difficult as it may be to accept, for the millions of investors who lamented after the bursting of the Technology Bubble in 2000, that "if I only got out earlier, I wouldn't have lost so much", they may have been correct.

As a result of the analysis of the classic stop-loss rule, I generalize the approach to general stopping rules. This generalization allows me to outline how stopping rules or superposition strategies can actually add or subtract value from underlying strategies, in the classic utility sense. I outline how uncertainty about return generating processes can produce situations, where various more complicated approximations of the true return generating process can lead to better investment strategies, by exploiting higher order characteristics of asset returns.

Following the general framework, I discuss various premiums related to the classic utility theory, such as, the stopping premium similar to my analysis of stop-loss rules. I then discuss the case of maximizing the stopping premium and examine how predictability in asset returns may create scenarios where stopping rules can add value to underlying strategies. To further demonstrate the performance impact of stopping rules I revisit the stop-loss rule as well as examine another popular rule, the buy-low sell-high strategy.

Appendix A

Supplemental Material

A.1 Asset Allocation

In this section, I present additional material related to asset allocation including basic definitions of investment options for retirement and tables which summarize the asset allocation specifications for lifecycle, target maturity, and balanced all-in-one asset allocation funds.

Basic Definitions: Retirement Investment Options

IRA: An Individual Retirement Account (IRA) is a brokerage account that allows earnings to compound over time on either a federally tax-free or tax-deferred basis. Investments in tax-advantaged accounts can compound more quickly than those in taxable accounts. Beyond an employer-sponsored retirement plan such as a 401(k), a Roth or Traditional IRA is widely considered the most advantageous retirement savings vehicle available.

401(k) Plan: This is the most popular of the defined contribution plans and is most commonly offered by larger employers. Employers often match employee contributions.

403(b) Tax-Sheltered Annuity Plan: Think of this as a 401(k) plan for employees of school systems and certain nonprofit organizations. Investments are

made in tax-sheltered annuities or mutual funds.

SIMPLE IRA: The Savings Incentive Match Plan for Employees of Small Employers is one of the newest types of employer-based retirement plans. There is also a 401(k) version of the SIMPLE.

Profit-Sharing Plan: The employer shares company profits with employees, usually based on the level of each employee's wages.

ESOP: Employee stock ownership plans are similar to profit-sharing plans, except that an ESOP must invest primarily in company stock. Under an ESOP, the employees share in the ownership of the company.

SEP: Simplified employee pension plans are used by both small employers and the self-employed. [1]

---

[1]For further references please consult: http://www.cfp.net http://www.fidelity.com http://www.vanguard.com http://www.tiaa-cref.org http://www.irs.org

Maturity	Equity (%)	Foreign Equity (%)	Bonds (%)	Investment Grade (%)	High Yield (%)	Short Term (%)
35	85	15	15	5	10	0
30	84	14	17	8	9	0
25	83	13	18	10	8	0
20	75	11	26	18	8	0
15	69	10	31	23	8	0
10	60	9	37	30	7	3
5	46	5	44	39	5	9
0	45	5	43	38	5	12
-5	23	0	38	37	1	40
<-5	20	0	40	39	1	40

Table A.1: Fidelity® Freedom Fund Allocation (Source: Fidelity Freedom® Funds Prospectus 2005, http://www.fidelity.com)

Maturity	Stocks (%)	Bonds (%)	Cash (%)
40	90	10	0
30	80	20	0
20	60	40	0
10	50	50	0
0	35	65	0
<0	20	75	5

Table A.2: Vanguard® Target Maturity Fund Allocation (Source: Vanguard® Target Maturity Funds Prospectus 2005, http://www.vanguard.com)

Maturity	Stocks (%)	Bonds (%)	Cash (%)
45	95	4	1
40	95	4	1
35	94	5	1
30	89	8	3
25	84	12	4
20	79	15	6
15	70	21	9
10	57	29	14
5	37	39	24
0	25	45	30

Table A.3: Putnam® RetirementReady Funds (Source: Putnam® RetirementReady Funds Prospectus 2005, http://www.putnam.com)

Maturity	Stocks (%)	Bonds (%)	Cash (%)
40	85	13	2
30	80	18	2
20	67	28	5
10	53	42	5
0	35	55	10

Table A.4: American Century® My Retirement Portfolios (Source: American Century® My Retirement Portfolios Prospectus 2005, http://www.americancentury.com)

Maturity	Stocks (%)	Bonds (%)	Cash (%)
40	93	7	0
35	93	7	0
30	93	7	0
25	93	7	0
20	87.5	12.5	0
15	80	19	1
10	73.5	23	3.5
5	67.5	25.5	7
0	60.5	30	9.5
<0	43	28.5	28.5

Table A.5: T. Rowe Price® Retirement Funds (Source: T. Rowe Price® Retirement Funds Prospectus 2005, http://www.troweprice.com)

Fund Name		w_1	w_2	t_1	t_2	α	β
Vanguard ®	Stocks	90	20	40	-5	0.015	0.28
	Bonds	10	75	40	-5	(-0.156,-0.105)	(0.72,0.697)
	Cash	0	5	5	-5	-0.005	0.025
Fidelity ®	Stocks	90	20	35	(15,5,0,-5,-10)	(0.0105,0.023,0.002,0.044,0.006)	(0.5325,0.345,0.45,0.45,0.26)
	Bonds	15	40	35	(15,5,0,-5,-10)	(-0.005,-0.014,0.004,0.012,-0.006)	(0.325,0.52,0.43,0.43,0.34)
	Cash	0	40	15	(5,0,-5,-10)	(-0.009,-0.006,-0.046,-0.01)	(0.135,0.12,0.12, 0.3)
T. Rowe Price ®	Stocks	90	20	25	-30	0.0127	(-0.127,-0.007,-0.127)
	Bonds	10	20	25	-30	(-0.127,-0.007,-0.127)	(0.418,0.345,0.218)
	Cash	0	20	20	-15	-0.0057	0.1143
Schwab ®	Stocks	85	20	40	-20	0.0163	0.525
	Bonds	13	50	40	-20	(-0.108,0.0007)	(0.5833,0.5073)
	Cash	0	30	5	-20	-0.0112	0.076
American Century ®	Stocks	85	35	40	(30,20,10,0)	(0.005,0.013,0.014,0.018)	(0.65,0.41,0.39,0.35)
	Bonds	13	55	40	(30,20,10,0)	(-0.005,-0.01,-0.014,-0.013)	(0.33,0.48,0.56,0.55)
	Cash	2	10	30	(20,10,0)	(-0.003,0,-0.005)	(0.11,0.05,0.1)

Table A.6: Summary of Linear Parameterizations for Target Maturity Portfolio Strategies over Stocks, Bonds, and Cash Equivalents for Vanguard®, Fidelity®, T. Rowe Price®, Schwab®, and American Century®(Source: http://www.vanguard.com http://www.fidelity.com http://www.troweprice.com http://www.schwab.com http://www.americancentury.com)

Vanguard LifeStrategy Funds ®	Stocks (%)	Bonds (%)
Income	20	80
Conservative Growth	40	60
Moderate Growth	60	40
Growth	80	20

Table A.7: Vanguard LifeStrategy®Funds: Performance Objective and Asset Allocation Specifications (Source: Vanguard LifeStrategy®Funds Prospectus 2005, http://www.vanguard.com)

A.2 Stop-Loss

In this appendix, I provide proofs of Propositions 1 and 2 in Sections A.2.1 and A.2.2, a derivation of the likelihood function of the behavioral regime-switching model (3.26) in Section A.2.3, and present some additional empirical results in Section A.2.4.

A.2.1 Proof of Proposition 1

The conclusion follows almost immediately from the observation that the conditional expectations in (3.4) and (3.6) are equal to the unconditional expectations because of the Random Walk Hypothesis (conditioning on past returns provides no incremental information), hence:

$$\Delta_\mu = -p_o\pi \leq 0 \qquad (A.1)$$

$$\frac{\Delta_\mu}{p_o} = -\pi \leq 0 \qquad (A.2)$$

and the other relations follow in a similar manner. ∎

Fidelity ®	Equity (%)	Foreign Equity (%)	Bonds (%)	Cash (%)	Horizon (years)
Most Aggressive	100	20	0	0	> 10
Aggressive Growth	85	15	15	0	> 10
Growth	70	10	25	5	> 5
Balanced	50	5	40	10	< 5
Conservative	50	0	20	30	≤ 5
Short-term	0	0	0	100	n/a

Table A.8: Fidelity® Funds: Performance Objective and Asset Allocation Specifications (Source: Fidelity® Funds Prospectus 2005, http://www.fidelity.com)

T. Rowe Price Personal Strategy Funds ®	Stocks (%)	Bonds (%)	Cash (%)
Income	40	40	20
Balanced	60	30	10
Growth	80	20	0

Table A.9: T. Rowe Price® Personal Strategy Funds: Performance Objective and Asset Allocation Specifications (Source: T. Rowe Price® Personal Strategy Funds Prospectus 2005, http://www.troweprice.com)

Schwab MarketTrack Portfolios ®	Stocks (%)	Bonds (%)	Cash (%)
Income	40	55	5
Balanced	60	35	5
Growth	85	15	5

Table A.10: Schwab® MarketTrack Portfolios:Performance Objective and Asset Allocation Specifications (Source: Schwab® MarketTrack Portfolios Prospectus 2005, http://www.schwab.com)

American Century One Choice Portfolios ®	Stocks (%)		Bonds (%)		Cash (%)	
Very Conservative	25	(20,30)	50	(45,55)	25	(15,35)
Conservative	45	(39,51)	45	(38,52)	10	(5,20)
Moderate	64	(53,73)	30	(21,41)	6	(0,15)
Aggressive	79	(60,90)	19	(10,30)	2	(0,15)
Very Aggressive	96	(75,100)	2	(0,10)	2	(0,15)

Table A.11: American Century® One Choice Portfolios: Performance Objective and Asset Allocation Specifications (Source: American Century® One Choice Portfolios Prospectus 2005, http://www.americancentury.com)

A.2.2 Proof of Proposition 2

Let r_t be a stationary AR(1) process:

$$r_t = \mu + \rho(r_{t-1} - \mu) + \epsilon_t \ , \quad \epsilon_t \overset{\text{IID}}{\sim} \text{White Noise}(0, \sigma_\epsilon^2) \ , \quad \rho \in (-1, 1] \tag{A.3}$$

We seek the conditional expectation of r_t given that the process is stopped out. If we let J be sufficiently large and $\delta = -\infty$, we note that $s_t = 0$ is equivalent to $R_{t-1}(J) < -\gamma$ and $s_{t-1} = 1$ with $R_{t-2}(J) \geq -\gamma$. Using log returns, we have

$$
\begin{aligned}
E[r_t | s_t = 0] &= E[r_t | R_{t-1}(J) < -\gamma, R_{t-2}(J) \geq -\gamma] &\text{(A.4)}\\
&= \mu(1-\rho) + \rho E[r_{t-1} + \epsilon_t | R_{t-1}(J) < -\gamma, R_{t-2}(J) \geq -\gamma] &\text{(A.5)}\\
&= \mu(1-\rho) + \rho E[r_{t-1} | R_{t-1}(J) < -\gamma, R_{t-2}(J) \geq -\gamma] &\text{(A.6)}
\end{aligned}
$$

By definition $R_{t-1}(J) \equiv r_{t-1} + \cdots + r_{t-J}$ and $R_{t-2}(J) = r_{t-2} + \cdots + r_{t-J-1}$. Setting $y \equiv r_{t-2} + \cdots + r_{t-J}$ then yields:

$$
\begin{aligned}
E[r_t | s_t = 0] &= \mu(1-\rho) + \rho E[r_{t-1} | R_{t-1}(J) < -\gamma, R_{t-2}(J) \geq -\gamma] &\text{(A.7)}\\
&= \mu(1-\rho) + \rho E_y\big[E[r_{t-1} | r_{t-1} < -\gamma - y, r_{t-J-1} \geq -\gamma - y]\big] &\text{(A.8)}
\end{aligned}
$$

125

For J large enough, the dependence between r_{t-J-1} and r_{t-1} is of order $\mathbf{o}(\rho^J) \approx 0$, hence:

$$\mathrm{E}_y\big[\mathrm{E}[r_{t-1}|r_{t-1} < -\gamma - y]\big] \leq \mathrm{E}_{r_{t-J-1}}\big[\mathrm{E}[r_{t-1}|r_{t-1} < r_{t-J-1}]\big] \tag{A.9}$$

$$\leq \mu - \sigma \tag{A.10}$$

which implies:

$$\mathrm{E}[r_t|s_t = 0] \leq \mu(1-\rho) + \rho(\mu - \sigma) \tag{A.11}$$

$$\leq \mu - \rho\sigma \tag{A.12}$$

∎

A.2.3   Behavioral Regime-Switching Likelihood Function

The behavioral regime-switching model begins with the standard regime-switching model (3.19):

$$r_t = I_t r_{1t} + (1 - I_t)r_{2t} \quad , \quad r_{it} \overset{\text{IID}}{\sim} \mathcal{N}(\mu_i, \sigma_i^2) \quad , \quad i = 1, 2$$

$$A \equiv \begin{array}{cc} & \begin{array}{cc} I_{t+1}=1 & I_{t+1}=0 \end{array} \\ \begin{array}{c} I_t=1 \\ I_t=0 \end{array} & \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} \end{array}$$

where I_t is an indicator function that takes on the value 1 when state 1 prevails and 0 when state 2 prevails, and A is the Markov transition probabilities matrix that governs the transitions between the two states.

The simple extension we propose is state-dependent transition probabilities:

$$\text{Prob}\,(I_t=0|I_{t-1}=1, \mathcal{F}_{t-1}; \theta) = \frac{\exp(a_1 + b_1 R_{t-1}(n))}{1 + \exp(a_1 + b_1 R_{t-1}(n))} \tag{A.13}$$

$$\text{Prob}\,(I_t=1|I_{t-1}=0, \mathcal{F}_{t-1}; \theta) = \frac{\exp(a_2 + b_2 R_{t-1}(n))}{1 + \exp(a_2 + b_2 R_{t-1}(n))} \tag{A.14}$$

where $R_{t-1}(n)$ is defined to be the cumulative n-period return:

$$R_{t-1}(n) = r_{t-1} + \cdots + r_{t-n} \qquad (A.15)$$

and $\mathcal{F}_{t-1}$ is the information set at time $t-1$, which includes r_{t-1}, $R_{t-1}(n)$, and all lags of these two variables.

Using methods from Hamilton (1994) we can construct the likelihood function as a function of the parameters $\theta \equiv \{\mu, \sigma, a_1, b_1, a_2, b_2\}$. Denote by $\mathbf{r}$ the matrix of data for equity and long-term bond returns from $t = 1, \ldots, T$. Then the likelihood function is given by:

$$f(\mathbf{r}|\theta) = \prod_{t=1}^{T} \Bigg(f(r_t|\mathcal{F}_{t-1}, I_t = 1; \theta)\text{Prob}(I_t = 1|\mathcal{F}_{t-1}; \theta) +$$

$$f(r_t|\mathcal{F}_{t-1}, I_t = 0; \theta)\text{Prob}(I_t = 0|\mathcal{F}_{t-1}; \theta) \Bigg) \qquad (A.16)$$

$$= \prod_{t=1}^{T} \Big(f(r_t|\mathcal{F}_{t-1}, I_t = 1; \theta)p_{1t} + f(r_t|\mathcal{F}_{t-1}, I_t = 0; \theta)p_{2t} \Big). \qquad (A.17)$$

The terms $f(r_t|\mathcal{F}_{t-1}, I_t = 1; \theta)$ and $f(r_t|\mathcal{F}_{t-1}, I_t = 0; \theta)$ are simply normal distributions for both bonds and equities. The conditional probabilities are more challenging. We present the expression for p_{1t} only, since the other conditional probability is similar:

$$p_{1t} \equiv \text{Prob}(I_t = 1|\mathcal{F}_{t-1}; \theta) = \text{Prob}(I_t = 1|I_{t-1} = 1, \mathcal{F}_{t-1}; \theta)q_{t-1}^1 +$$

$$\text{Prob}(I_t = 1|I_{t-1} = 0, \mathcal{F}_{t-1}; \theta)q_{t-1}^2) \qquad (A.18a)$$

$$= \left(1 - \frac{\exp(a_1 + b_1 R_{t-1}(n))}{1 + \exp(a_1 + b_1 R_{t-1}(n))}\right)q_{1t-1} + \frac{\exp(a_2 + b_2 R_{t-1}(n))}{1 + \exp(a_2 + b_2 R_{t-1}(n))}q_{2t-1} \qquad (A.18b)$$

$$= (1 - g_1(R_{t-1}(n)))q_{1t-1} + g_2(R_{t-1}(n))q_{2t-1} \qquad (A.18c)$$

where

$$q_{1t-1} \equiv \frac{f(\mathcal{F}_{t-1}|I_{t-1} = 1, \mathcal{F}_{t-2}; \theta)p_{1t-2}}{f(\mathcal{F}_{t-1}|I_{t-1} = 1, \mathcal{F}_{t-2}; \theta)p_{1t-2} + f(\mathcal{F}_{t-1}|I_{t-1} = 0, \mathcal{F}_{t-2}; \theta)p_{2t-2}} \qquad (A.19a)$$

$$q_{2t-1} \equiv \frac{f(\mathcal{F}_{t-1}|I_{t-1} = 0, \mathcal{F}_{t-2}; \theta)p_{2t-2}}{f(\mathcal{F}_{t-1}|I_{t-1} = 1, \mathcal{F}_{t-2}; \theta)p_{1t-2} + f(\mathcal{F}_{t-1}|I_{t-1} = 0, \mathcal{F}_{t-2}; \theta)p_{2t-2}} \qquad (A.19b)$$

We are left with one final term which we must characterize, $f(\mathcal{F}_{t-1}|I_{t-1}=1, \mathcal{F}_{t-2}; \theta)$, which is the probability density function for the new information set given the past information and the past state. Since $\mathcal{F}_{t-1} = \{(r_{t-1}, r_{ft-1}), \mathcal{F}_{t-2}\}$ we need the same expression $f(r_{t-1}|I_{t-1}=1, \mathcal{F}_{t-2}; \theta)$ which is a normal distribution.

Denote by $\phi(\cdot)$ the standard normal density function, and let:

$$\phi_{it} \equiv \phi\left(\frac{r_t - \mu_i}{\sigma_i}\right) \qquad i = 1, 2.$$ (A.20)

Then the likelihood function may be rewritten more compactly as:

$$f(\mathbf{r}|\theta) = \prod_{t=1}^{T} (\phi_{1t}p_{1t} + \phi_{2t}p_{2t}) , \quad \text{where}$$ (A.21a)

$$p_{1t} = (1 - g_1(R_{t-1}))\frac{\phi_{1t-1}p_{1t-1}}{\phi_{1t-1}p_{1t-1} + \phi_{2t-1}p_{2t-1}} +$$

$$g_2(R_{t-1})\frac{\phi_{2t-1}p_{2t-1}}{\phi_{1t-1}p_{1t-1} + \phi_{2t-1}p_{2t-1}}$$ (A.21b)

$$p_{2t} = g_1(R_{t-1})\frac{\phi_{1t-1}p_{1t-1}}{\phi_{1t-1}p_{1t-1} + \phi_{2t-1}p_{2t-1}} +$$

$$(1 - g_2(R_{t-1}))\frac{\phi_{2t-1}p_{2t-1}}{\phi_{1t-1}p_{1t-1} + \phi_{2t-1}p_{2t-1}}$$ (A.21c)

We can then use an iterative algorithm that calculates p_{it} as a function of R_{t-1}, r_{t-1}, and p_{it-1}. Once we have all the p_{it}'s, we substitute them into the expression for $f(\mathbf{r}|\theta)$ to calculate the likelihood function for a given θ, and then solve for the maximum likelihood estimator in the usual fashion. ∎

A.2.4 Additional Empirical Results

In this section, we provide four additional tables to supplement the empirical results in the main text. In Table A.12, we present a more detailed set of summary statistics for the buy-and-hold equities strategy of Section 3.5 with and without the stop-loss policy, including means, standard deviations, Sharpe ratios, and skewness and kurtosis coefficients for various stop-loss parameters (γ, δ, J). In Tables A.13 and A.14, we present similar performance statistics, but only for returns from the stopped-out

periods, assuming a re-entry threshold of 0% in Table A.13 and 2% in Table A.14. And in Table A.15, we report p-values of Kolmogorov-Smirnov test statistics designed to distinguish between the unconditional returns of our two asset classes and their conditional counterparts, conditioned on being stopped-out.

J (δ=0)	γ (%)	μ (%)	σ (%)	Sharpe	Skew	Kurt	μ (%)	σ (%)	Sharpe	Skew	Kurt
No Stops	—	12.5	14.4	0.87	-0.3	4.7	12.5	14.4	0.87	-0.3	4.7
		δ = 0					δ = 2%				
3	-4	12.9	13.2	0.97	-0.5	5.2	12.1	13.0	0.94	-0.4	5.3
	-6	12.8	13.5	0.95	-0.4	5.0	12.6	13.4	0.94	-0.4	5.0
	-8	13.2	13.7	0.96	-0.4	4.9	13.1	13.5	0.97	-0.4	5.0
	-10	12.8	13.8	0.93	-0.4	4.8	12.8	13.7	0.93	-0.4	4.8
	-12	12.7	13.9	0.92	-0.4	4.7	12.7	13.8	0.92	-0.4	4.7
	-14	12.5	14.0	0.89	-0.4	4.7	12.5	13.9	0.90	-0.4	4.7
6	-4	13.5	13.2	1.03	-0.5	5.3	12.8	12.8	1.00	-0.5	5.5
	-6	13.3	13.4	1.00	-0.5	5.0	12.8	13.1	0.97	-0.5	5.2
	-8	13.2	13.5	0.98	-0.5	5.0	12.8	13.3	0.97	-0.5	5.0
	-10	13.1	13.6	0.96	-0.4	4.9	12.9	13.4	0.96	-0.4	5.0
	-12	12.7	13.7	0.93	-0.4	4.8	12.5	13.5	0.92	-0.4	4.9
	-14	12.5	13.7	0.91	-0.4	4.8	12.3	13.6	0.90	-0.4	4.8
12	-4	13.7	13.4	1.03	-0.5	5.1	13.5	13.0	1.03	-0.5	5.3
	-6	13.6	13.5	1.01	-0.5	5.0	13.4	13.1	1.02	-0.5	5.2
	-8	13.4	13.5	0.99	-0.5	4.9	13.1	13.3	0.98	-0.5	5.1
	-10	13.3	13.6	0.98	-0.5	4.9	13.0	13.3	0.97	-0.5	5.0
	-12	13.1	13.6	0.96	-0.5	4.9	12.8	13.4	0.95	-0.5	4.9
	-14	13.0	13.8	0.95	-0.4	4.8	12.9	13.5	0.95	-0.5	4.9
18	-4	13.1	13.6	0.96	-0.5	4.9	12.9	13.4	0.96	-0.5	4.9
	-6	13.2	13.6	0.97	-0.5	4.8	13.0	13.5	0.97	-0.5	4.9
	-8	13.3	13.7	0.98	-0.5	4.8	13.1	13.5	0.97	-0.5	4.8
	-10	13.6	13.7	0.99	-0.5	4.8	13.4	13.6	0.99	-0.5	4.8
	-12	13.7	13.7	1.00	-0.5	4.8	13.4	13.6	0.98	-0.5	4.8
	-14	13.6	13.7	0.99	-0.5	4.8	13.6	13.6	1.00	-0.4	4.9

Table A.12: Performance statistics of a buy-and-hold strategy for the CRSP Value-Weighted Total Market return with and without a simple stop-loss-policy, where the stop-loss asset yields the Ibbotson Associates Long-Term Government Bond return, for stop-loss thresholds $\gamma = 4$–14%, re-entry threshold $\delta = 0\%$, 2%, and window sizes $J = 3$, 6, 12, and 18 months, from January 1950 to December 2004.

J (δ=0) Months	γ (%)	μ (%)	μ_s (%)	σ (%)	σ_s (%)	Sharpe	Sharpe_s	Skew	Skew_s	Kurt	Kurt_s	P_<0	P_<0,s	P_c|s	P_c|s,s	Stop-Loss Statistics			
																#Stops Per Year	Total Stops	#Months Per Year	Total Months
3	-4	11.6	14.7	20.1	9.2	0.58	1.60	0.1	0.5	2.8	2.8	0.42	0.41	0.52	0.57	0.89	49	1.25	69
	-6	14.6	18.0	20.6	9.9	0.71	1.83	-0.1	0.3	2.9	2.5	0.38	0.38	0.48	0.50	0.58	32	0.91	50
	-8	11.4	23.1	22.2	9.4	0.51	2.45	0.0	0.0	2.6	2.4	0.40	0.29	0.51	0.46	0.40	22	0.64	35
	-10	20.9	28.4	24.6	10.4	0.85	2.72	-0.2	-0.3	2.4	2.3	0.33	0.25	0.50	0.42	0.29	16	0.44	24
	-12	18.6	25.2	26.5	9.4	0.70	2.67	0.0	-0.7	2.3	2.3	0.35	0.24	0.53	0.35	0.20	11	0.31	17
	-14	26.1	20.5	31.5	9.5	0.83	2.16	-0.2	-0.3	2.0	1.8	0.30	0.30	0.40	0.50	0.13	7	0.18	10
6	-4	5.2	13.3	19.1	9.0	0.27	1.47	0.3	0.2	3.1	3.1	0.47	0.38	0.59	0.52	1.00	55	1.44	79
	-6	5.4	13.8	19.7	9.6	0.27	1.44	0.3	0.3	3.2	2.8	0.46	0.41	0.59	0.54	0.75	41	1.15	63
	-8	2.5	13.0	22.4	10.6	0.11	1.22	0.4	0.3	2.7	2.6	0.52	0.43	0.64	0.52	0.49	27	0.80	44
	-10	-0.8	10.7	24.5	10.4	-0.03	1.03	0.4	0.2	2.4	2.7	0.52	0.39	0.58	0.52	0.36	20	0.56	31
	-12	8.8	12.8	25.3	10.4	0.35	1.24	0.3	0.2	2.3	2.8	0.48	0.40	0.60	0.48	0.31	17	0.45	25
	-14	10.7	8.0	25.5	9.2	0.42	0.87	0.3	-0.1	2.3	2.2	0.48	0.43	0.48	0.48	0.27	15	0.42	23
12	-4	2.6	14.9	19.8	9.5	0.13	1.58	0.4	0.2	3.1	2.8	0.53	0.35	0.61	0.55	0.95	52	1.13	62
	-6	3.3	16.1	20.3	9.9	0.16	1.62	0.4	0.1	3.1	2.6	0.53	0.36	0.62	0.53	0.84	46	1.00	55
	-8	4.5	15.6	20.4	10.1	0.22	1.56	0.5	0.1	3.2	2.6	0.54	0.36	0.64	0.52	0.67	37	0.91	50
	-10	3.3	15.4	21.7	10.5	0.15	1.46	0.5	0.1	2.9	2.5	0.53	0.37	0.63	0.53	0.51	28	0.78	43
	-12	5.9	15.2	22.8	10.2	0.26	1.48	0.4	-0.1	2.8	2.5	0.50	0.33	0.61	0.50	0.38	21	0.65	36
	-14	2.0	14.3	24.0	10.6	0.08	1.35	0.5	-0.4	2.7	2.0	0.50	0.38	0.62	0.46	0.27	15	0.47	26
18	-4	8.1	16.1	20.6	10.5	0.39	1.53	0.4	0.1	3.1	2.5	0.49	0.34	0.57	0.53	0.73	40	0.85	47
	-6	5.2	15.5	20.9	10.6	0.25	1.46	0.5	0.1	3.2	2.5	0.52	0.34	0.59	0.55	0.67	37	0.80	44
	-8	3.3	17.0	21.8	11.0	0.15	1.54	0.5	0.1	3.1	2.4	0.53	0.34	0.61	0.53	0.58	32	0.69	38
	-10	-3.3	18.0	23.2	11.5	-0.14	1.56	0.7	0.0	3.1	2.3	0.63	0.31	0.69	0.47	0.44	24	0.58	32
	-12	-5.3	19.8	24.0	11.2	-0.22	1.77	0.8	-0.1	3.1	2.5	0.66	0.28	0.72	0.45	0.36	20	0.53	29
	-14	-10.2	17.2	23.6	10.4	-0.43	1.65	0.8	-0.2	3.3	2.8	0.65	0.27	0.73	0.46	0.25	14	0.47	26

Table A.13: Performance statistics during stopped-out periods of a buy-and-hold strategy for the CRSP Value-Weighted Total Market return with and without a simple stop-loss-policy, where the stop-loss asset yields the Ibbotson Associates Long-Term Government Bond return, for stop-loss thresholds $\gamma = 4$-14%, re-entry threshold $\delta = 0\%$, and window sizes $J = 3, 6, 12,$ and 18 months, from January 1950 to December 2004. The subscript S denotes performance in the stop-loss asset: Ibbotson Associate's Long-Term Government Bond return.

J (δ=2%) Months	γ(%)	μ (%)	μ_s (%)	σ (%)	σ_s (%)	Sharpe	Sharpe_s	Skew	Skew_s	Kurt	Kurt_s	P_{<0}	P_{<0,s}	P_{cjs}	P_{cjs,s}	Stop-Loss Statistics			
																#Stops Per Year	Total Stops	#Months Per Year	Total Months
3	-4	13.3	10.4	19.1	9.1	0.69	1.14	0.0	0.4	2.9	3.1	0.41	0.43	0.47	0.56	0.82	45	1.69	93
	-6	12.9	13.7	20.1	9.9	0.64	1.39	-0.1	0.3	2.9	2.7	0.42	0.40	0.50	0.52	0.53	29	1.13	62
	-8	9.0	17.5	22.0	10.0	0.41	1.76	0.0	0.0	2.5	2.5	0.44	0.33	0.53	0.51	0.35	19	0.78	43
	-10	18.6	25.0	25.1	11.1	0.74	2.26	-0.3	-0.4	2.3	2.4	0.33	0.26	0.44	0.41	0.27	15	0.49	27
	-12	15.9	21.2	26.8	10.4	0.59	2.03	-0.1	-0.7	2.2	2.4	0.35	0.25	0.45	0.40	0.18	10	0.36	20
	-14	18.8	16.5	32.0	11.1	0.59	1.49	-0.2	-0.5	1.8	2.1	0.33	0.33	0.42	0.42	0.11	6	0.22	12
6	-4	10.3	11.7	18.9	8.7	0.55	1.34	0.2	0.2	3.0	3.2	0.45	0.37	0.55	0.52	0.87	48	1.85	102
	-6	10.1	11.8	19.8	9.4	0.51	1.26	0.1	0.3	3.0	2.9	0.44	0.40	0.53	0.55	0.67	37	1.40	77
	-8	7.4	11.1	22.7	10.6	0.33	1.05	0.2	0.3	2.5	2.6	0.48	0.42	0.56	0.54	0.44	24	0.95	52
	-10	3.7	10.0	25.0	10.6	0.15	0.95	0.2	0.1	2.2	2.6	0.47	0.39	0.56	0.50	0.35	19	0.65	36
	-12	12.7	11.6	25.6	10.6	0.49	1.10	0.1	0.1	2.1	2.7	0.43	0.40	0.53	0.50	0.29	16	0.55	30
	-14	14.5	7.6	25.8	9.6	0.56	0.79	0.0	-0.2	2.2	2.2	0.43	0.43	0.43	0.50	0.25	14	0.51	28
12	-4	4.6	12.3	20.2	9.4	0.23	1.30	0.3	0.2	2.8	2.9	0.51	0.37	0.58	0.54	0.71	39	1.42	78
	-6	5.3	12.9	20.5	9.8	0.26	1.31	0.3	0.1	2.8	2.7	0.51	0.38	0.61	0.54	0.60	33	1.29	71
	-8	7.5	12.7	20.6	10.0	0.37	1.27	0.3	0.1	2.9	2.7	0.51	0.37	0.62	0.52	0.49	27	1.18	65
	-10	7.4	12.9	22.3	10.4	0.33	1.24	0.3	0.1	2.6	2.6	0.53	0.37	0.61	0.51	0.40	22	0.93	51
	-12	8.8	12.2	23.4	10.4	0.38	1.17	0.3	-0.1	2.5	2.4	0.50	0.36	0.57	0.48	0.29	16	0.76	42
	-14	3.3	9.2	23.7	10.4	0.14	0.89	0.4	-0.2	2.6	2.0	0.54	0.40	0.60	0.46	0.22	12	0.64	35
18	-4	8.4	12.0	20.8	11.5	0.40	1.04	0.3	-0.4	2.9	3.6	0.49	0.34	0.56	0.46	0.53	29	1.11	61
	-6	4.7	10.3	20.9	11.7	0.23	0.89	0.4	-0.4	3.0	3.6	0.53	0.37	0.58	0.49	0.47	26	1.04	57
	-8	3.9	10.8	21.7	12.2	0.18	0.88	0.4	-0.4	3.4	3.4	0.52	0.38	0.58	0.48	0.42	23	0.91	50
	-10	-0.7	14.5	22.9	12.9	-0.03	1.12	0.5	-0.6	2.9	3.5	0.59	0.31	0.64	0.44	0.35	19	0.71	39
	-12	-0.3	14.4	23.3	12.9	-0.01	1.11	0.5	-0.7	2.8	3.6	0.59	0.30	0.65	0.43	0.31	17	0.67	37
	-14	-7.3	15.7	24.2	10.6	-0.30	1.48	0.6	-0.3	2.8	2.8	0.63	0.27	0.70	0.47	0.20	11	0.55	30

Table A.14: Performance statistics during stopped-out periods of a buy-and-hold strategy for the CRSP Value-Weighted Total Market return with and without a simple stop-loss-policy, where the stop-loss asset yields the Ibbotson Associates Long-Term Government Bond return, for stop-loss thresholds $\gamma = 4$–14%, re-entry threshold $\delta = 2\%$, and window sizes $J = 3, 6, 12,$ and 18 months, from January 1950 to December 2004. The subscript S denotes performance in the stop-loss asset: Ibbotson Associate's Long-Term Government Bond return.

γ (%)	p-Value r_e	p-Value r_b	p-Value r_e	p-Value r_b	p-Value r_e	p-Value r_b	p-Value r_e	p-Value r_b
	J = 2		J = 6		J = 12		J = 18	
-3	0.08	0.12	0.04	0.04	0.20	0.11	0.05	0.02
-4	0.04	0.02	0.07	0.05	0.07	0.10	0.09	0.07
-5	0.09	0.01	0.18	0.01	0.10	0.09	0.08	0.10
-6	0.08	0.01	0.18	0.10	0.10	0.05	0.04	0.04
-7	0.12	0.00	0.07	0.24	0.04	0.12	0.08	0.04
-8	0.04	0.00	0.02	0.13	0.08	0.08	0.08	0.08
-9	0.02	0.00	0.01	0.07	0.11	0.04	0.05	0.09
-10	0.00	0.00	0.00	0.17	0.06	0.05	0.01	0.03
-11	0.00	0.00	0.01	0.12	0.12	0.01	0.01	0.03
-12	0.01	0.00	0.05	0.10	0.06	0.02	0.01	0.05
-13	0.09	0.00	0.03	0.20	0.02	0.02	0.01	0.04
-14	0.04	0.06	0.07	0.25	0.02	0.02	0.01	0.04

Table A.15: p-values of Kolmogorov-Smirnov tests for for the equality of the empirical distributions of monthly returns unconditionally and after stop-loss triggers, for the CRSP Value-Weighted Total Market Index and Ibbotson Associates Long-Term Bond Index from January 1950 to December 2004.

A.3 Stochastic Dominance

When solving the investor's problem, it is necessary to specify both a preference structure for investors, as well as parameterize the underlying distribution of asset returns. It can be argued that making a particular assumption, about either preferences or underlying asset price dynamics, will dramatically influence the results. To circumvent this problem, it is possible to analyze conditional asset returns using first and second order stochastic dominance. The key advantages of stochastic dominance is that it is applicable for general preferences and it deals directly with empirical data without distributional assumptions. Although the approach is non-parametric, it is definitely not a panacea for distributional assumptions. In fact, the application of stochastic dominance testing has been limited due to computational issues, questions about the impact of sampling error on results for small samples, as well as limited model flexibility to account for multiple asset choices.

Empirical tests for stochastic dominance can be divided into two main approaches: grid based methods (Davidson and Duclous 2000), and Kologorov-Smirnov based methods (McFadden 1989, Kaur, Rao, and Singh 1994, Barrett and Donald 2003). Davidson and Duclous (2000) propose a grid method which can be useful for dealing with multiple hypotheses allowing for both dependent and independent samples. In the following sections, I summarize basic definitions for first and second order stochastic dominance and discuss a summary of two methods for empirical testing of stochastic dominance. For the empirical testing, I focus on the Kolmogorov-Smirnov approach; this is due to its simplicity and the possibility to use conditional bootstraping methods to account for possible dependence in samples. The Kolmogorov-Smirnov based approach was first outlined by McFadden (1989), but various authors have extended his results. In particular, Barrett and Donald (2003) propose a bootstraping based method, which focuses on two prospects and allows both for dependent sampling with unequal sample lengths.

A.3.1 Basic Definitions

In this section, I summarize, according to Huang and Litzenberger (1988), various types of first and second order stochastic dominance.

First Order Stochastic Dominance

Following Huang and Litzenberger (1988), a risky asset A is said to **first order stochastically dominate** B, or $A \leq_1 B$ if all individuals having monotonically increasing and continuous utility functions would prefer A to B. Letting $\tilde{r}_A$ and $\tilde{r}_B$ be the rate of return on the assets A and B respectively. $F(\cdot)$ is the cumulative distribution function. The following three statements are equivalent:

$$1. \quad A \leq_1 B \tag{A.22}$$

$$2. \quad F_A(z) \leq F_B(z) \tag{A.23}$$

$$3. \quad \tilde{r}_A = \tilde{r}_B + \tilde{\alpha} \quad \tilde{\alpha} \geq 0 \tag{A.24}$$

Second Order Stochastic Dominance

According to Huang and Litzenberger (1988), a risky asset A is said to **second order stochastically dominate** B, denoted $A \leq_2 B$ if all risk averse individuals having utility functions with continuous derivatives except on a countable subset of $[1,2]$ prefer A to B. $A \leq_2 B$ if and only if

$$E[\tilde{r}_A] = E[\tilde{r}_B] \tag{A.25}$$

$$\text{and } S(y) = \int_c^d (F_A(z) - F_B(z))dz \leq 0 \quad \forall y \in [c,d] \tag{A.26}$$

then again the following three statements are equivalent

$$1. \quad A \leq_2 B \tag{A.27}$$

$$2. \quad E[\tilde{r}_A] = E[\tilde{r}_B] \text{ and } S(y) \leq 0 \quad \forall y \tag{A.28}$$

$$3. \quad \tilde{r}_A =^d \tilde{r}_B + \tilde{\epsilon}, \text{ with } E[\tilde{\epsilon}|\tilde{r}_A] \tag{A.29}$$

Monotonic Second Order Stochastic Dominance

A risky asset A second order stochastic monotonically dominates B, denoted $A \leq_2^M B$ if all risk averse and non-satiable individuals prefer A to B, then the following three expressions are equivalent.

1. $A \leq_2^M B$ (A.30)

2. $E[\tilde{r}_A] \geq E[\tilde{r}_B]$ and $S(y) \leq 0$ $\forall y$ (A.31)

3. $\tilde{r}_A =^d \tilde{r}_B + \tilde{\epsilon}$, with $E[\tilde{\epsilon}|\tilde{r}_A]$ (A.32)

A.3.2   Empirical Tests for Stochastic Dominance

In this section, I summarize two methods for performing empirical tests for stochastic dominance. Let $r_i^e, i = 1 \ldots N$ be N independent and identically distributed (iid) sample returns with cumulative distribution function (CDF), $F_e(x)$, and $r_i^b, i = 1 \ldots N$ be iid samples returns with CDF, $F_b(x)$, define $D_e(x)$ and $D_b(x)$ as follows:

$$D_e(x) = \int_0^x F_e(u)du$$

$$D_b(x) = \int_0^x F_b(u)du$$

Without loss of generality, I assume the support is $[0, \bar{x}]$, then e can be said to stochastically dominates b of the second order if $D^e(x) \leq D^b(x), \forall x \geq 0$. There are two alternative hypotheses H_0 and H_1.

$$H_0 : D_e(x) \leq D_b(x) \quad \forall x \in [0, \bar{x}]$$

$$H_1 : D_e(x) > D_b(x) \quad for \; some \; x \in [0, \bar{x}]$$

Barrett and Donald (2003) propose the following test statistic, $\hat{K}$ to test H_0.

$$\hat{K} = \left(\frac{N^2}{2N}\right)^{1/2} \sup_x [\hat{D}_e(x) - \hat{D}_b(x)]$$

$$where\; D_e(x) = \int_0^x (x - e_i)dF_w(u)$$

$$= \frac{1}{N}\sum_{i=1}^{N}(x - e_i)I(e_i \leq x)$$

$$= \frac{1}{N}\sum_{i=1}^{N}(x - e_i)_+$$

By adopting asymptotic results for Brownian Bridges, Barrett and Donald (2003) show that the test will reject with probability one if the hypothesis H_0 is false. To test this hypothesis, p-values can be simulated using an arbitrarily fine grid for calculating the suprema of $\hat{K}$(see Barrett and Donald 2003).

To deal with higher order stochastic dominance, Barrett and Donald (2003) suggest other test statistics which are derived using asymptotic results as well. Let $\{X\}_{i=1}^{N}$ be N samples from a marginal distribution $F_X(\cdot)$ and $\{Y_i\}_{i=1}^{M}$ be M samples from a marginal distribution $G_Y(\cdot)$, they define the following statistics $\hat{S}_1$ and $\hat{S}_2$ for first and second order stochastic dominance respectively.

$$\hat{S}_1 = \left(\frac{NM}{N+M}\right)^{1/2} \sup_z (\hat{G}(z) - \hat{F}(z))$$

$$\hat{S}_2 = \left(\frac{NM}{N+M}\right)^{1/2} \sup_z (I(\hat{G}(z)) - I(\hat{F}(z)))$$

$$where\; I(\hat{F}(z)) = \frac{1}{N}\sum_{i=1}^{N}(z - x_i)_+$$

$$I(\hat{G}(z)) = \frac{1}{M}\sum_{i=1}^{M}(z - y_i)_+$$

The test for the hypothesis that G first order and second order stochastically dominates F can be summarized by the following two sets of hypothesises: H_0^1 and H_1^1,

and H_0^2 and H_1^2.

$$H_0^1: \qquad G(x) \le F(x) \quad \forall x$$

$$H_1^1: \qquad G(x) > F(x) \quad for \ some \ x$$

$$H_0^1: \qquad \int^x G(t)dt \le \int^x F(t)dt \quad \forall x$$

$$H_1^1: \qquad \int^x G(t)dt > \int^x F(t)dt \quad for \ some \ x$$

Using the fact that the limiting distributions for both F and G can be characterized by Brownian Bridge Processes combined with the continuous mapping theorem, Barrett and Donald (2003) define the random variables $\bar{S}_1^{G,F}$ and $\bar{S}_2^{G,F}$.

$$\bar{S}_1^{G,F} = \sup_z (\sqrt{\lambda} B \cdot G - \sqrt{1-\lambda} B \cdot F)$$

$$\bar{S}_2^{G,F} = \sup_z (\sqrt{\lambda} \int B \cdot G dt - \sqrt{1-\lambda} \int B \cdot F dt)$$

Where $B \cdot F$ is the composite of a brownian motion and the distribution function F, and λ is the sampling frequency for both distributions. The corresponding probability of rejecting the null hypothesis for both first and second order stochastic dominance can be bounded above by $\bar{S}$ using Proposition 1 from Barrett and Donald (2003).

$$\lim_{N,M \to \infty} P(reject H_0) \quad \le \quad P(\bar{S}^{F,G} \le c) = \alpha(c)$$

The corresponding p-values can be evaluated using simulation, to calculate the probability that $\bar{S} > \hat{S}$ for both first and second order stochastic dominance. By taking R resamples of size (N, M) of the empirical distribution, and letting $\hat{F}_{N,r}^*$ and $\hat{G}_{M,r}^*$ be the corresponding empirical distributions of the resamples. An estimate of the pvalue, denoted $\hat{p}_1$ and $\hat{p}_2$, can be computed using Monte Carlo simulation over the

R samples.

$$\hat{p}_1 \approx \frac{1}{R} \sum_{r=1}^{R} 1(\bar{S}_1^{F,G} > \hat{S}_1)$$

$$\approx \frac{1}{R} \sum_{r=1}^{R} 1(\sqrt{\left(\frac{NM}{N+M}\right)} \sup_x (\hat{G}_{M,r}^* - \hat{F}_{N,r}^*) > \hat{S}_1)$$

$$\hat{p}_2 \approx \frac{1}{R} \sum_{r=1}^{R} 1(\bar{S}_2^{F,G} > \hat{S}_2)$$

$$\approx \frac{1}{R} \sum_{r=1}^{R} 1(\sqrt{\left(\frac{NM}{N+M}\right)} \sup_x (\int^x \hat{G}_{M,r}^*(t)dt - \int^x \hat{F}_{N,r}^*(t)dt) > \hat{S}_1)$$

Agnew, J., P. Balduzzi, and A. Sunden, 2003, "Portfolio Choice and Trading in a Large 401(K) Plan", *American Economic Review* 93, 193–215.

Agnew, J., 2003, "An Analysis of How Individuals React to Market Returns in One 401(K) Plan", Working paper, The College of William and Mary.

Ait-Sahalia, Y. and M. Brandt, 2001, "Variable Selection for Portfolio Choice", *Journal of Finance* 56, 1297-1351.

Alexander, S., 1961, "Price Movements in Speculative Markets: Trends or Random Walks", *Industrial Management Review* 2, 7–26.

Ameriks, John, and Stephen Zeldes, 2004, "How Do Household Portfolio Shares Vary with Age?", working paper, Columbia University.

Ang, A. and G. Bekaert, 2004, "How Regimes Affect Asset Allocation", *Financial Analysts Journal* 60, 86–99.

Ang, A., Bekaert, G., and J. Liu, 2005, "Why Stocks May Disappoint", *Journal of Financial Economics* 76, 471–508.

Balduzzi, P., and A.W. Lynch, 1999, "Transaction Costs and Predictability: Some Utility Cost Calculations," *Journal of Financial Economics* 52, 47-78.

Barber, B. and T. Odean, 2000, "Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors", *Journal of Finance* 55, 773-806.

Barber, B. and T. Odean, 2005, "All that Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors", unpublished working paper.

Barberis, N., 2000, " Investing in the Long Run when Returns are Predictable," *Journal of Finance* 55, 225-264.

Barberis, N., and Thaler, R., "Chapter 18: A Survey of Behavioral Finance" in Handbook of the Economics and Finance, edited by G.M. Constantinides, M. Harris, and R. Stulz, 2003 Elsevier Science B.V.

Barrett, G. and S. Donald, 2003, "Consistent Tests for Stochastic Dominance", *Econometrica* 71, 71–104.

Bell, D., 1982a, "Regret in Decision Making under Uncertainty", *Operations Research* 30, 961–981.

Bell, D., 1982b, "Risk Premiums for Decision Regret", *Management Science* 29, 1156–1166.

Bell, D., 1985, "Disappointment in Decision Making under Uncertainty", *Operations Research* 33, 1–27.

Benartzi, S., and R. Thaler, 2001, " Naive Diversification Strategies in Defined Contribution Savings Plans", *American Economic Review* 91, 79-98.

Benartzi, S. and R. Thaler, 2001, "Myopic loss aversion and the equity premium puzzle", *Quarterly Journal of Economics* 110, 73-92.

Berkelaar, A., and R. Kouwenberg, 2000, "Dynamic Asset Allocation and Downside-Risk Aversion", working paper.

Berkelaar, A.,Kouwenberg, R., and Post, T., 2003, "Optimal Portfolio Choice under Loss Aversion", working paper.

Biais, B., P. Hillion, and C. Spatt, 1995, An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse", *Journal of Finance* 50, 1655–1689.

Blume, M., and I. Friend, 1978, *The Changing Role of the Individual Investor.* New York: John Wiley.

Bollerslev, T., 1986, "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31, 307–327.

Brennan, M. and Y. Xia, 2001, "Persistence, Predictability, and Portfolio Planning", working paper.

Brandt, M., 2004, "Portfolio Choice Problems", in Y. Ait-Sahalia and L.P. Hansen (eds.), *Handbook of Financial Econometrics*, forthcoming. New York: Elsevier.

Brandt, M., 1999, "Estimating Portfolio and Consumption Choice," *Journal of Finance* 54, 1609-1645.

Brandt, M., Goyal, A., Santa-Clara, P., and J. Stroud, 2001, " A simulation approach to dynamic portfolio choice with an application to learning about return predictability", *Review of Economics and Statistics*, forthcoming.

Brandt, M., Goyal, A., Santa-Clara, P., and J. Stroud, 2005, "A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning About Return Predictability", *Review of Financial Studies* 18, 831–873.

Brennan, M., Schwartz, E., and R. Lagnado, 1997, "Strategic Asset Allocation", *Journal of Economic Dynamics and Control* 21, 1377-1403.

Brennan, M. and Y. Xia, 2001, "Persistence, Predictability, and Portfolio Planning", unpublished working paper.

Calvet, L., Campbell, J., and P., Sodini, 2006, "Down or Out: Assesing the Welfare Costs of Household Investment Mistakes", Working Paper No. 12030, NBER.

Campbell, J., Chan, Y., and L. Viceira, 2003, "A Multivariate Model of Strategic Asset Allocation", *Journal of Financial Economics* 67, 41–80.

Campbell, J., Lo, A., and C. MacKinlay, 1997, *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.

Campbell, J., and L. Viceira, 1999, " Consumption and Portfolio Decisions when Expected Returns are Time-Varying", *Quarterly Journal of Economics* 114, 433–495.

Campbell, J., and L. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long Term Investors*, Oxford University Press: New York, NY.

Campbell, J., Lo, A., and C. MacKinlay, 1997, *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.

Carr, P., and R. Jarrow, 1990, "The Stop-Loss Start-Gain Paradox and Option Valuation: A New Decomposition into Intrinsic and Time Value", *Review of Financial Studies* 3, 469–492.

Chakravarty, S., and C. Holden, 1995, "An Integrated Model of Market and Limit Orders", *Journal of Financial Intermediation* 4, 213–241.

Chang, F., 1988, "The Inverse Optimal Problem: A Dynamic Programming Approach", *Econometrica* 56, 147–172.

Clarke, R., Krase, S., and M. Statman, 1994, "Tracking Errors, Regret, and Tactical Asset Allocation", *Journal of Portfolio Management* 20, 16–24.

Constantinides, G. 1976, "Stochastic Cash Management with Fixed and Proportional Transaction Costs", *Management Science* 22, 1320–1331.

Constantinides, G., 1983, " Capital Market Equilibrium with Personal Tax", *Econometrica* 51, 611-636.

Constantinides, G., 1984, " Optimal Stock Trading with Personal Taxes: Implications for Prices and the Abnormal January Returns", *Journal of Financial Economics* 13, 65-89.

Constantinides, G., 1986. "Capital Market Equilibrium with Transaction Costs", *Journal of Political Economy* 94, 842–862.

Constantinides, G., and M., Magill, 1976, "Portfolio Selection with Transaction Costs", *Journal of Economic Theory* 13, 264–271.

142

Cox, J. and C. Huang, 1989, "Optimum Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process", *Journal of Economic Theory* 49, 33–83.

Das, S., and R. Sundaram, 2000, "A Numerical Algorithm for Optimal Consumption-Investment Problems", unpublished working paper, Harvard University.

Davidson, R., and J. Duclous, 2000, "Statistical Interference for Stochastic Dominance and for the Measurement Of Poverty And Inequality", *Econometrica* 68, 1435–1464.

DeBondt, W. and R. Thaler, 1986, "Does the Stock Market Overreact?", *Journal of Finance* 40, 793–807.

Easley, D., and M. O'Hara, 1991, "Order Form and Information in Securities Markets", *Journal of Finance* 46, 905–927.

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

Fama, E., 1965, "The Behavior of Stock Market Prices", *Journal of Business* 38, 34–105.

Fama, E. and M. Blume, 1966, "Filter Fules and Stock Market Trading Profits", *Journal of Business* 39, 226–241.

Fama, E. and K. French, 1988, "Permanent And Temporary Components of Stock Prices", *Journal of Political Economy* 96, 246–273.

Fischoff, B. and P. Slovic, 1980, "A Little Learning...: Confidence in Multicue Judgment Tasks", in R. Nickerson (ed.), *Attention and Performance, VIII*. Hillsdale, NJ: Erlbaum.

Fong, W., Wong, W., and H.H. Lean, 2003, "International Momentum Strategies: A Stochastic Dominance Approach", unpublished working paper, National University of Singapore.

Genesove, D., and Mayer, C., 2001, "Loss Aversion and Seller Behavior: Evidence from the Housing Market", NBER Working Paper No. W8143.

Gervais, S. and T. Odean, 2001, "Learning to Be Overconfident", *Review of Financial Studies* 14, 1–27.

Gilboa, I., and D. Schmeidler, 1989, "Maximum Expected Utility with Non-Unique Prior", *Journal of Mathematcial Economics* 18, 141–153.

Goetzmann, W. and M. Massa, 2003, "Index Funds and Stock Market Growth", *Journal of Business* 76, pp. 1–27.

Goetzmann, W., Massa, M. and G. Rouwenhurst, 1999, "Behavioral Factors in Mutual Fund Flows", Yale International Center for Finance, Working Paper No. 00–14, Yale University.

Grinblatt, M., and M, Keloharju, 2001, "What Makes Investors Trade?", *Journal of Finance* 56(2), 589-616.

Grinblatt, M, and Han, B., 2001, "The Disposition Effect and Momentum", unpublished working paper, University of California, Los Angeles.

Grinblatt, M, and Han, B., 2004, "Prospect Theory, Mental Accounting, and Momentum", to appear in *Journal of Financial Economics.*

Gul, F., 1991, "A Theory of Disappointment Aversion", *Econometrica* 59, 667–686.

Hamilton, J., 1989, "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", *Econometrica* 57, 357–384.

Hamilton, J., 1994, *Time Series Analysis*, Princeton University Press.

Handa, P., and R. Schwartz, 1996, "Limit Order Trading", *Journal of Finance* 51, 1835–1861.

Harris, L. and J. Hasbrouck, 1996, "Market vs. Limit Orders: The SuperDot Evidence on Order Submission Strategy", *Journal of Financial and Quantitative Analysis* 31, 213–231.

Heaton, J., and D. Lucas, 2000, " Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk", *Journal of Finance* 55, 1163-1198.

Huang, C., and R. Litzenberger, 1988, *Foundations for Financial Economics*. New Jersey: Prentice Hall.

Huberman, G. and T. Regev, 2001, "Contagious Speculation and a Cure for Cancer: A Nonevent That Made Stock Prices Soar", *Journal of Finance* 56, 387–396.

Ingersoll, J., 1987, *Theory of Financial Decision Making*. Totowa, NJ: Rowman & Littlefield.

Jegadeesh, N., 1990, "Evidence of Predictable Behavior of Security Returns", *Journal of Finance* 45, 881–898.

Jegadeesh, N., and S. Titman, 1993, "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency", *Journal of Finance* 48, 65–91.

Kahneman, D., and A. Tversky, 1979, "Prospect Theory: An Analysis of Decision Under Risk", *Econometrica* 47, 263–291.

Kahneman, D., and Tversky, A., 1992, "Advances in Prospect Theory: Cumulative Representation of Uncertainty", *Journal of Risk and Uncertainty* 5, 297–323.

Kaminski, K., 2006, "General Superposition Strategies", unpublished manuscript.

Kaur, B., Rao, S. and H. Singh, 1994, "Testing of Second-Order Stochastic Dominance of Two Distributions", *Econometric Theory* 10, 849–866.

Kelly, M., 1995, "All their eggs in one basket: Portfolio diversification of US households", *NBER Working Paper* 9685.

Kim, T., and E. Omberg, 1996, "Dynamic Non-Myopic Portfolio Behavior", *Review of Financial Studies* 9, 141–161.

Kogan, L., and R. Uppal, 2000, "Risk Aversion and Optimal Portfolio Policies in Partial and General Equilibrium Economies", unpublished working paper, University of Pennsylvania.

Laibson, D., 1997, "Golden Eggs and Hyperbolic Discounting", *Quarterly Journal of Economics* 62, 443–477.

Lichtenstein, S., Fischoff, B. and L. Phillips, 1982, "Calibration of Probabilities: The State of the Art to 1980", in D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases.* Cambridge, UK: Cambridge University Press.

Lintner, J., 1965, "The Valuation of Risk Assets and the Selection of Risky Investments In Stock Portfolios and Capital Budgets", *Review of Economics and Statistics* 47, 13–37.

Liu, J., 1999", Portfolio Selection in Stochastic Environments", unpublished working paper, University of California, Los Angeles.

Liu, J. 2001, "Dynamic Portfolio Choice and Risk Aversion", working paper, 2001.

Lo, A., 1991, "Long-Term Memory in Stock Market Prices", *Econometrica* 59, 1279–1313.

Lo, A., 2002, "The Statistics of Sharpe Ratios", *Financial Analysts Journal* 58, 36–50.

Lo, A. and C. MacKinlay, 1988, "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test", *Review of Financial Studies* 1, 41–66.

Lo, A. and C. MacKinlay, 1990, "An Econometric Analysis of Nonsynchronous Trading", *Journal of Econometrics* 45, 181–212.

Lo, A. and C. MacKinlay, 1999, *A Non-Random Walk Down Wall Street.* Princeton, NJ: Princeton University Press.

Lo, A., MacKinlay, C. and J. Zhang, 2002, "Econometric Models of Limit-Order Executions", *Journal of Financial Economics* 65, 31–71.

Lo, A., Mamaysky, H., and J. Wang, 2004, "Asset Prices and Trading Volume under Fixed Transactions Costs", *Journal of Political Economy* 112, 1054-1090.

Lo, A. and J. Wang, 1995, "Implementing Option Pricing Models When Asset Returns are Predictable", *Journal of Finance* 50, 87–129.

Loomes, G., and Sugden, R., 1982, "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty", *The Economic Journal* 92, 805–824.

Lucas, R., 1978, "Asset Prices in an Exchange Economy", *Econometrica* 46, 1429–1446.

Markowitz, H., 1952, "Portfolio Selection", *Journal of Finance* 7, 77–91.

Massa, M., and A. Simonov, 2004, "Hedging, Familiarity and Portfolio Choice", CEPR Discussion Paper No. 4789.

McFadden, D., 1989, "Testing for Stochastic Dominance, " in T.B. Fomby and T.K. Seo (eds.), *Studies in the Economics of Uncertainty.* New York: Springer-Verlag.

Merton, R. 1969, "Lifetime Portfolio Selection Under Uncertainty; The Continuous-Time Case", *Review of Economics & Statistics* 51, 247–257.

Merton, R., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model", *Journal of Economic Theory* 3, 373–413.

Merton, R., 1973, "An Intertemporal Capital Asset Pricing Model", *Econometrica* 41, 867–887.

Nofsinger, J., and Sias, R., 1999, " Herding and Feedback Trading by Institutional and Individual Investors ", *Journal of Finance* 54(6).

Odean, T., 1998, "Are Investors Reluctant to Realize Their Losses?", *Journal of Finance* 53, 1775–1798.

Odean, T., 1999, "Do Investors Trade Too Much", *American Economic Review* 89, 1279–1298.

Polkovnichenko, V., 2002, "Household Portfolio Diversification", unpublished working paper, University of Minnesota.

Poterba, J. and L. Summers, 1988, "Mean Reversion in Stock Returns: Evidence and Implications", *Journal of Financial Economics* 22, 27–60.

Quiggin, J., 1982, "A theory of anticipated utility", *Journal of Economic Behavior and Organization* 3, 323-343.

Ranguelova, E., 2001, " Disposition Effect and Firm Size: New Evidence in Individual Trading Activity", working paper, Harvard University.

Samuelson, P. 1969, "Lifetime Portfolio Selection By Dynamic Stochastic Programming", *Review of Economics and Statistics* 51, 239–246.

Seppi, D., 1997, "Liquidity Provision with Limit Orders and a Strategic Specialist", *Review of Financial Studies* 10, 103–150.

Segal, U., 1987, " Some remarks on Quiggin's anticipated utility", *Journal of Economic Behavior and Organization* 8, 145-154.

Segal, U., 1989, " Anticipated utility: a measure representation approach", *Annals of Operations Research* 19, 359-373.

Seyhun, H.N., 1993, " Can omitted risk factors explain the January effect? A stochastic dominance approach," *Journal of Financial and Quantitative Analysis*, 28, 195-212.

Shapira, Z., and Venezia, I., 2001, "Patterns of Behavior of Professionally Managed and Independent Investors," *Journal of Banking and Finance* 25, 1573-1587.

Sharpe, W., 1964, "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk", *Journal of Finance* 19, 425–442.

Sharpe, W., 1994, "The Sharpe Ratio", *Journal of Portfolio Management* Fall, 49–58.

Shefrin, H., 2005, *Behavioral Approaches to Asset Pricing*. New York: Elsevier.

Shefrin, M. and M. Statman, 1985, "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence", *Journal of Finance* 40, 777–790.

Shefrin, M., and M. Statman, 1994, "Behavioral Capital Asset Pricing Theory", *Journal of Financial and Quantitative Analysis* 29, 323-349.

Shiller, R., 2001, "Human Behavior and the Efficiency of the Financial System", *Cowles Foundation for Research in Economics* Paper No. 1025, 1305–1340.

Spurgin, R., 2001, "How to Game Your Sharpe Ratio", *The Journal of Alternative Investments* 4, 38–46.

Stigler, G., *The Theory of Price*, 3d Ed. New York: Macmillan, 1966.

Thaler, R., 1980, " Toward a Positive Theory of Consumer Choice", *Journal of Economic Behavior and Organization* 13, 253-282.

Thaler, R., 1985, "Mental Accounting and Consumer Choice", *Marketing Science* 4, 199-214.

Thaler, R., 1994, "Psychology and Savings Policies", *American Economic Review* 84, 186–192.

Thaler, R., 1999, " Mental Accounting Matters", *Journal of Behavioral Decision Making* 12(3),183-206.

Thaler, R., and Johnson, E., 1990, "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes in Risky Choice", *Management Science* 36(6), 643-660.

Thaler, R. and H. Shefrin, 1981, "An Economic Theory of Self Control", *The Journal of the Political Economy* 89,392–406.

Tobin, J., 1958, "Liquidity Preference as Behavior Towards Risk", *Review of Economic Studies* 25, 65–86.

Tracy, J., Schneider, H., and S. Chan, 1999, " Are Stocks Overtaking Real Estate in Household Portfolios?", *Current Issues in Economics and Finance* 5:5, Federal Reserve Bank of New York.

Tsay, R., *Analysis of Financial Time Series*, Wiley, 2002.

Tschoegl, A., 1988, "The Source and Consequences of Stop Orders: A Conjecture", *Managerial and Decision Economics* 9, 83–85.

Tversky, A. and D. Kahneman, 1981, "The Framing of Decisions and the Psychology of Choice", *Science* 211, 453–458.

Wachter, J., 2002, "Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets", *Journal of Financial and Quantitative Analysis* 37, 63–91.

Xia, Y., 2001, "Learning about Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation", *Journal of Finance* 56, 205–246.

Yaari, M., 1987, " The dual theory of choice under risk", *Econometrica* 55,95-115.