Dynamic Trading and Behavioral Finance

by

Alexander Remorov

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 2016

Author . . . .
Sloan School of Management
April 26, 2016

Certified by . . . .
Andrew W. Lo
Charles E. and Susan T. Harris Professor
Director of Laboratory for Financial Engineering
Thesis Supervisor

Accepted by . . . .
Dimitris Bertsimas
Boeing Leaders for Global Operations Professor
Co-Director, Operations Research Center

Dynamic Trading and Behavioral Finance

by

Alexander Remorov

Submitted to the Sloan School of Management
on April 26, 2016, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

The problem of investing over time remains an important open question, considering the recent large moves in the markets, such as the Financial Crisis of 2008, the subsequent rally in equities, and the decline in commodities over the past two years. We study this problem from three aspects.

The first aspect lies in analyzing a particular dynamic strategy, called the stop-loss strategy. We derive closed-form expressions for the strategy returns while accounting for serial correlation and transactions costs. When applied to a large sample of individual U.S. stocks, we show that tight stop-loss strategies tend to underperform the buy-and-hold policy due to excessive trading costs. Outperformance is possible for stocks with sufficiently high serial correlation in returns. Certain strategies succeed at reducing downside risk, but not substantially. We also look at optimizing the stop-loss level for a class of these strategies.

The second approach is more behavioral in nature and aims to elicit how various market players *expect to react* to large changes in asset prices. We use a global survey of individual investors, financial advisors, and institutional investors to do this. We find that most institutional investors expect to exhibit highly contrarian reactions to past returns in terms of their equity allocations. Financial advisors are also mostly contrarian; a few of them demonstrate passive behavior. In contrast, individual investors are, on average, extrapolative, and can be partitioned into four distinct types: passive investors, risk avoiders, extrapolators, and everyone else.

The third part of the thesis studies how people *actually* trade. We propose a new model of dynamic trading in which an investor is affected by behavioral heuristics, and carry out extensive simulations to understand how the heuristics affect portfolio performance. We propose an MCMC algorithm that is reasonably successful at estimating model parameters from simulated data, and look at the predictive ability of the model. We also provide preliminary results from looking at trading data obtained from a brokerage firm. We focus on understanding how people trade their portfolios conditional on past returns at various horizons, as well as on past trading behavior.

Thesis Supervisor: Andrew W. Lo
Title: Charles E. and Susan T. Harris Professor

Director of Laboratory for Financial Engineering

Acknowledgments

I would like to thank my advisor, professor Andrew Lo, for his continuous guidance, support, and motivation throughout my thesis. I am very grateful to have been working with him on very interesting and exciting projects that I truly enjoyed. Andrew's brilliant intuition and big picture thinking, wealth of knowledge and new ideas, and amazing work ethic made it a great honor and pleasure to work with him and learn from him. His advice and thoughts on various aspects of academia, industry and life in general have been and definitely will be extremely helpful for me throughout my entire career.

I am very grateful for continuous research discussion and guidance with my other committee members: professors Dimitris Bertsimas, Leonid Kogan, and Hui Chen. I learned a lot about the power of analytics and optimization from the classes I took with Dimitris and conversations I had with him. Leonid and Hui were extremely helpful in teaching me and providing advice on how to build and extend financial models, especially from an asset pricing perspective in which they are truly experts.

It was great to work with Daniel Elkind during the last few months of my degree. The project on modeling investor trading really took off after he came on board, and we have been able to produce great results at a very fast pace due to his great enthusiasm and hard work, as well as very productive discussion. I wish him all the best during the remainder of his Ph.D. and I am sure he will do great.

The research collaboration that Andrew and I had with Natixis Asset Management was crucial for producing the results of this thesis. David Goodsell and Stephanie Giardina did a great job in designing and carrying out investor surveys; it was also very helpful to continuously discuss the progress and the results with them. Support from John Hailer is also very much appreciated.

I also want to acknowledge the terrific help of the U.S. retail brokerage firm that provided us with the trading data towards the end of my thesis. The individuals at the firm responsible for delivering the data were extremely responsive and involved in compiling and sending the data and ensuring it is clean and satisfies our joint research purposes. Even though only preliminary results of this joint project are included in this thesis, I am truly excited about this project and hope that it will produce some amazing results down the road.

I am grateful to Laura Rose and Andrew Carvalho at the ORC who were always very helpful with the many questions I had throughout my time here and took care of any relevant administrative items in a very timely and helpful manner. I also want to thank Andrew's secretaries Jayna Cummings, Patsy Thompson, and Allison McDonough for being very responsive and always being able to fit me into Andrew's busy schedule.

My Ph.D. would not be the same without all the great people I met here. Charles T., Velibor M., Andrew L., Will M., Shalev B., Anna P., Stefano T., Alexandre S., Daniel S., Ali A., Alex W., Virgile G., Joey H., Zach O., Arthur F., Adam E., David F., Ludovica R., Mariapola T., and Dan S. are some of the awesome students that made the last four years such a fun and fulfilling experience. It was also an awesome experience leading the ORC soccer team, and I want to thank my fellow teammates Nikita, Anurag, Kevin, Marnix, Adam, Osman, Fransisco, Elliott, Joey, Virgile, Max, Hai, Alex S., Alex W., David, Will, and Charles, among others for coming to the pick-up and intramurals games and playing the great game together. I have a lot of fond memories from the two years I lived with my two roommates and very close friends Danny Shi and David Rush, especially when we played Settlers (and particularly the games which I won).

I want to thank Edwin Cao for being an amazing friend. We had a lot of entertaining conversations and heated discussions on all kinds of topics ranging from analytics to Texas to nuggets, and I truly cherish the fact that we were able to talk almost every day besides often being in different cities or even countries. Edwin's continuous grind to be tier 1 or nothing will always continue to be a great source of motivation for me.

I want to thank my family for their continuous encouragement and support throughout my thesis and also throughout my whole life. Particularly my mother and father, who have taught me so much and have always done their best for me. I would not be here without them.

Finally, I want to thank my wife Kristina for all her care and love. For being so understanding and supportive. For being a constant source of inspiration and happiness. For always being by my side. I love you.

Contents

Chapter 1

Stop-loss Strategies with Serial Correlation, Regime Switching and Transaction Costs

(joint work with Andrew W. Lo)

Abstract

Stop-loss strategies are commonly used by investors to reduce their holdings in risky assets if prices or total wealth breach certain pre-specified thresholds. We derive closed-form expressions for the impact of stop-loss strategies on asset returns that are serially correlated, regime switching, and subject to transaction costs. When applied to a large sample of individual U.S. stocks, we show that tight stop-loss strategies tend to underperform the buy-and-hold policy in a mean-variance framework due to excessive trading costs. Outperformance is possible for stocks with sufficiently high serial correlation in returns. Certain strategies succeed at reducing downside risk, but not substantially. We also look at optimizing the stop-loss level for a class of these strategies.

Keywords: Stop-loss strategy; Risk management; Investments; Portfolio management; Asset allocation; Behavioral finance

JEL classification: G11, G12

1.1 Introduction

Many investors attempt to limit the downside risk of their investments by using stop-loss strategies, the most common of which is the stop-loss order, a standing order to liquidate a position when a security's price crosses a pre-specified threshold. By closing out the position, the investor is hoping to avoid further losses.

If prices follow random walks, any price movement in the past has no bearing on future returns—as long as the risky asset has a positive risk premium, the investor's portfolio will have a higher expected return by staying invested in the asset rather than liquidating it after its price reaches a particular limit. In this case, Kaminski and Lo (2014) have shown that the stop-loss strategy tends to underperform a buy-and-hold strategy. However, there is extensive evidence that financial asset prices do not follow random walks (e.g., Lo and MacKinlay, 1988; Poterba and Summers, 1988; Jegadeesh and Titman, 1993). A natural question is whether these departures from randomness can be exploited using a dynamic investment strategy, including stop-loss policies.

In this paper, we focus on simple dynamic strategies incorporating stop-loss rules to determine how they compare to static buy-and-hold strategies. We provide closed-form expressions for the returns of a large class of these strategies and derive conditions under which they underperform or outperform buy-and-hold. Assuming that prices follow a first-order autoregressive process, we prove that the log-returns of "tight" stop-loss strategies—strategies with price triggers that are close to the asset's current price—are approximately linear in the interaction term between autocorrelation and volatility, providing an explicit relation between the profitability of a stop-loss policy, return predictability, and volatility. This expression yields bounds on how large return autocorrelation and volatility must be to beat a buy-and-hold strategy after accounting for trading costs.

We also consider the dynamic optimization problem of an investor trading a risky asset and a risk-free asset, with the returns of the risky asset following an AR(1) process. Consistent with our approximation results, optimal stop-loss behavior arises only when we have positive serial correlation in returns; furthermore, an investor should use a tighter stop-loss level for higher values of serial correlation and volatility of the process.

We extend our theoretical analysis by simulating various return processes and by comparing the performance of stop-loss and buy-and-hold policies in a mean-variance framework. We consider two general processes—an AR(1) and a regime-switching process—and vary the underlying parameters for each. In the first case, with a high enough serial correlation and volatility, the stop-loss strategy provides superior risk-adjusted returns in comparison to the buy-and-hold strategy. In the regime-switching case, the stop-loss strategy gives better performance in a few cases, and this outperformance comes from a large reduction in volatility rather than an improvement in raw returns. We also look at the tail performance of the strategy, as measured by skewness and maximum drawdown. We find that if a longer horizon for past returns is used to make the decision whether to stop out or not, downside risk tends to improve over the buy-and-hold.

To illustrate the practical relevance of stop-loss strategies, we perform a detailed empirical analysis of the performance of these strategies using a large sample of U.S. stock returns from 1964 to 2014. To derive realistic measures of performance, we incorporate transaction costs in our backtests by using bid-ask spreads, as well as historical estimates when such spreads are missing.[1] Our empirical findings are most relevant to short-term traders, who usually employ tight stop-loss policies and frequently change their positions. We find that the performance of tight stop-loss strategies is closely related to the realized return autocorrelation over the investment period, which supports the common trading adage: “The trend is your friend." However, such strategies require a lot of trading, leading to high transaction costs. As a result, tight stop-loss strategies are able to outperform the buy-and-hold strategies only when asset returns are significantly serially correlated.

Of course, a stop-loss rule alone does not fully define an entire investment strategy since, after exiting a risky investment, the investor must decide when to re-enter. We consider several simple re-entry policies as part of our definition of a stop-loss rule and demonstrate that it is usually beneficial to re-invest soon after being stopped out in the case of tight stops. Another aspect that must be considered is where cash is invested after a stop-out. Assuming that cash is immediately invested in a risk-free asset, we show that the risk-free rate has a significant impact on the effectiveness of a stop-loss strategy, and this impact reconciles some of the inconsistencies among existing empirical studies of stop-loss strategies.

From a broader perspective, the use of stop-loss strategies can correct for the tendency of investors to hold on to losers too long and to sell winners too early, a behavioral

[1] For transaction costs prior to 1993, we use the Hasbrouck (2009) dataset.

bias known as the “disposition effect” first documented by Shefrin and Statman (1985). The presence of behavioral biases such as this has been well documented in the finance literature.[2] While most of this research has focused on the empirical evidence for these biases and the theoretical models to explain them, few studies have proposed methods for investors to actively avoid or protect against such biases. Stop-loss strategies are an important first step in this direction.

1.2 Literature review

Kaminski and Lo (2014) lay out the first general framework for analyzing stop-loss strategies. They start with analytical results for the performance of a stop-loss policy and consider three cases for the return process of the risky asset. For a simple random walk, the policy always produces lower expected returns. For an AR(1) process, the policy improves performance in the case of momentum, but hurts performance in the case of mean reversion. For a two-state Markov regime-switching model, the strategy sometimes gives better performance, since it tends to outperform the buy-and-hold strategy only in the low-mean state.

There are a few other analytical studies of stop-loss strategies. Glynn and Iglehart (1995) derive an optimal strategy by demonstrating that the expected value of the stock price at the time of exit satisfies a relatively simple ordinary differential equation (ODE). They also present an example of a utility function with a very heavy penalty on losses, which would lead the investor to set up a finite stop-loss limit. This contrasts with the case of constant relative risk aversion (CRRA) utility, where it is optimal to not use a stop-loss (Merton, 1969). Glynn and Iglehart's ODE approach was later applied to derive the optimal selling rule in more complicated settings for the return distribution, including for a regime-switching process (Zhang, 2001; Pemy, 2011) and a mean-reverting process (Zhang and Zhang , 2008; Ekström, Lindberg, and Tysk, 2011). Besides the ODE approach, Abramov, Khan, and Khan (2008) analyze the trailing stop strategy in a discrete time framework, while Esipov and Vaysburd (1999) present a partial differential equation approach for analyzing stop-loss policies.

With respect to the empirical literature on stop-loss strategies, Kaminski and Lo

[2]For surveys of behavioral biases, see Hirshleifer (2001) and Shefrin (2010).

(2014) consider the strategy of investing in the S&P 500, using monthly frequency for the historical returns and U.S. long-term bonds as the “safe” asset. Erdestam and Stangenberg (2008) and Snorrason and Yusupov (2009) study the strategies applied to stocks in the OMX Stockholm 30 Index. Lei and Li (2009) use daily data for individual U.S. stocks (with the S&P 500 and the one-month U.S. T-bill considered as the “safe" assets) and conclude that traditional stop-loss strategies are able to reduce losses for some stocks, but not for others. Trailing stop-loss strategies are found to consistently reduce investment risk.

There has been little attention paid to the transaction costs associated with stop-loss strategies. Two papers that do address this issue are Macrae (2005) and Detko, Ma, and Morita (2008). They note that in many cases, the associated hidden costs, such as slippage, result in lower strategy returns.

Finally, stop-loss strategies may also benefit investors by implicitly correcting for some of their behavioral biases. One such bias is the disposition effect, where investors tend to hold losers for too long and sell winners too early, as documented by Shefrin and Statman (1985), Ferris, Haugen and Makhija (1988), and Odean (1998). Wong, Carducci, and White (2006) find evidence for the disposition effect in an experimental setting and propose using stop-loss orders to offset this bias. However, studies exploring this possibility have yielded mixed and inconclusive results. For example, Garvey and Murphy (2004) investigate a sample of trading records for professional traders in the U.S. and find that, while traders tend to use stop-loss orders and avoid large losses, they still exhibit the disposition effect. Nevertheless, Richards, Rutterford, and Fenton-O’Creevy (2011) find that retail investors who employ stop-loss strategies exhibit the disposition effect to a smaller extent than those who do not.

We build on the existing literature in several important directions. The first is an advanced theoretical formula approximating stop-loss strategy performance when returns follow an AR(1) process. We link the conclusions from this formula to historical stop-loss performance. The second direction is that we rigorously incorporate transaction costs into our analysis of simulated and historical strategy performance. The sample of assets we consider is individual U.S. stocks, which has different dynamics to the S&P 500 futures considered in the Kaminski and Lo (2014) paper. The third piece is that we consider downside risk by looking at skewness and maximum drawdown as performance metrics.

Finally, we perform extensive simulations to gain insights on how the performance of the strategies is related to their specifications and the parameters of the underlying returns process.

1.3 Analytical results

We consider the stop-loss strategy introduced by Kaminski and Lo (2014) and later generalized by Erdestam and Stangenberg (2008). We invest 100% in the risky asset at the start of the period. If its cumulative return over J consecutive periods drops below a specified threshold γ, we liquidate our position and invest in the risk-free asset; otherwise we stay fully invested. To buy the asset again, the cumulative return over I periods has to exceed a threshold δ.

Denote by r_t the log return on the risky asset at time t. Define the cumulative log return $R_t(N)$ over N consecutive periods as:

$$R_t(N) \equiv \sum_{j=1}^{N} r_{t-j+1} \, . \tag{1}$$

Let s_t be the proportion of wealth allocated to the risky asset at the start of period t. We define the stop-loss strategy as:

Definition 1. *A* **fixed rolling-window policy** $\mathcal{S}(\gamma, \delta, J, I)$ *is a dynamic asset allocation rule* $\{s_t\}$ *between the risky asset* Q *and the safe asset* F*, such that:*

$$s_t = \begin{cases} 1 \text{ if } R_{t-1}(J) > \log(1+\gamma) \text{ and } s_{t-1} = 1 \text{ (stay in)} \\ 0 \text{ if } R_{t-1}(J) \leq \log(1+\gamma) \text{ and } s_{t-1} = 1 \text{ (exit)} \\ 0 \text{ if } R_{t-1}(I) < \log(1+\delta) \text{ and } s_{t-1} = 0 \text{ (stay out)} \\ 1 \text{ if } R_{t-1}(I) \geq \log(1+\delta) \text{ and } s_{t-1} = 0 \text{ (re-enter)} \end{cases} \tag{2}$$

The strategy can be implemented in practice as follows. During each day we track the log cumulative return over the past J days, where J is specified by the strategy. The return over the current day is also included in the calculation. As we approach the close of the day, if the cumulative return drops below the specified threshold, we sell the asset. We thus assume that the asset price does not move significantly just prior to the close.

Since sometimes selling the asset right before the close may be problematic, we propose a modified stop-loss strategy that is more realistic to implement. At the start of each day, if the cumulative return over the previous J days (not counting the current day) is below a specified threshold, we submit a market-on-close order to sell the asset at the end of the day. We call this strategy the *delayed fixed rolling window policy* $\mathcal{S}^d(\gamma, \delta, J, I)$; the formal definition of this policy is given in the Appendix.

We next present a theoretical analysis of the performance of the stop-loss strategy when the underlying returns follow an AR(1) process. We give an explicit expression for the returns of the strategy after accounting for the dependence on the risk-free rate, as well as transaction costs. This enables us to analyze when the stop-loss beats the buy-and-hold strategy in terms of raw returns.

1.3.1 Strategy returns for an AR(1) process

Suppose we are investing over a period of length T and hold the risky asset on the first day. The return on the safe asset is assumed to be constant and equal to r_f in each period, while trading costs (as a percentage of capital) are also assumed to be constant at c per period in which a transaction on the risky asset is made.[3]

The log returns $\{r_t\}$ on the risky asset follow an AR(1) process:

$$r_t \;=\; \mu \;+\; \rho(r_{t-1} - \mu) \;+\; \epsilon_t, \;\; \epsilon_t \sim WN(0, \sigma^2), \tag{3}$$

where $\rho \in (-1, 1)$ is a constant.

Consider the stop-loss strategy $\mathcal{S}(\gamma, \delta, J, I)$. We restrict ourselves to cases where γ and δ are small, while $J = I = 1$. This corresponds to a tight stop-loss/start-gain strategy in which we exit or re-enter the risky asset if its one-day return is too low or too high, respectively.

Let $a \equiv \log(1 + \gamma)$, $b \equiv \log(1 + \delta)$. In proposition 1 we present an approximation for the performance of the strategy absent any trading costs:

Proposition 1. *Assume $|\rho|$ is not too large. If $b \geq a$, the expected log-return of the*

[3] Thus c includes both the cost of one transaction on the risky asset and one transaction on the risk-free one, since whenever we sell one of the assets, we immediately invest in the other.

stop-loss strategy $\mathcal{S}(\gamma, \delta, 1, 1)$ is approximated by:

$$
\begin{aligned}
\mathbb{E}[R_{sp}] \;\approx\;& \pi\left[1+(T-1)\Big(\Phi(\frac{\mu-b}{\tilde\sigma})+p_1\big(\Phi(\frac{b-\mu}{\tilde\sigma})-\Phi(\frac{a-\mu}{\tilde\sigma})\big)\Big)\right] \;+\; Tr_f \;+ \\
& \rho\frac{\tilde\sigma}{\sqrt{2\pi}}(T-1)\left[\exp(-\frac{(\mu-b)^2}{2\tilde\sigma^2})+p_1\big(\exp(-\frac{(a-\mu)^2}{2\tilde\sigma^2})-\exp(-\frac{(b-\mu)^2}{2\tilde\sigma^2})\big)\right],
\end{aligned}
\tag{4}
$$

where:

$$
p_1 \;=\; \frac{\mathbb{P}(r_{t-1}\ge b,\; a<r_t<b)}{\mathbb{P}(r_{t-1}\ge b,\; a<r_t<b)+\mathbb{P}(r_{t-1}\le a,\; a<r_t<b)} \quad,\quad \tilde\sigma^2 \;\equiv\; \frac{\sigma^2}{1-\rho^2}\;.
$$

If $b<a$, then the expected return is approximately:

$$
\begin{aligned}
\mathbb{E}[R_{sp}] \;\approx\;& \pi\left[1+(T-1)(\Phi(\frac{\mu-a}{\tilde\sigma})+p_2(\Phi(\frac{a-\mu}{\tilde\sigma})-\Phi(\frac{b-\mu}{\tilde\sigma})))\right] \;+\; Tr_f \;+ \\
& \rho\frac{\tilde\sigma}{\sqrt{2\pi}}(T-1)\left[\exp(-\frac{(\mu-a)^2}{2\tilde\sigma^2})+p_2(\exp(-\frac{(b-\mu)^2}{2\tilde\sigma^2})-\exp(-\frac{(a-\mu)^2}{2\tilde\sigma^2}))\right],
\end{aligned}
\tag{5}
$$

where:

$$
p_2 \;=\; \frac{\mathbb{P}(r_{t-1}\le a,\; b\le r_t\le a)}{\mathbb{P}(r_{t-1}\le a,\; b\le r_t\le a)+\mathbb{P}(r_{t-1}\ge b,\; b\le r_t\le a)} \quad,\quad \tilde\sigma^2 \;\equiv\; \frac{\sigma^2}{1-\rho^2}\;.
$$

The first part in (4) is the return contributed from the mean μ and is similar to the random walk case. However, in this case we have another part that depends on the autocorrelation coefficient ρ and volatility σ. In fact, for small values of $|\rho|$, $\tilde\sigma$ does not depend too much on ρ and as a result the second part of (4) is close to linear in $\rho\sigma$.

To get an idea of how much the serial correlation adds to the return, we consider the case when $\mu \approx 0$ and $a = b = 0$. Here, we use a very tight stop-loss/start-gain strategy and assume low daily returns on the risky asset. We then have:

$$
\mathbb{E}(R_{sp}) \;\approx\; \pi(1+\frac{1}{2}(T-1)) \;+\; \rho\frac{\tilde\sigma}{\sqrt{2\pi}}(T-1) \;+\; Tr_f. \tag{6}
$$

Assuming as before a risk-free rate of 0, and no trading costs, in order to beat the expected log-return of the buy-and-hold strategy, we need to have:

$$
\pi(1+\frac{1}{2}(T-1)) \;+\; \rho\frac{\tilde\sigma}{\sqrt{2\pi}}(T-1) \;+\; Tr_f \;>\; \mu T \;\Leftrightarrow\; \rho > \frac{\pi\sqrt{2\pi}}{2\tilde\sigma}\;. \tag{7}
$$

Assuming ρ is small, we have $\tilde{\sigma} \approx \sigma$, and as a result, an approximate lower bound on ρ is:

$$\rho \; > \; \frac{\sqrt{2\pi}}{2}\frac{\pi}{\sigma} \; \approx \; 1.25\frac{\pi}{\sigma} \; \approx \; 1.25\frac{\mu}{\sigma}, \tag{8}$$

in order to beat the buy-and-hold strategy. For daily U.S. stock data over the 1964–2014 period, the ratio of daily return to standard deviation is 5.62%, implying that on average a serial correlation of around 7.0% or higher is necessary to beat the buy-and-hold strategy.

1.3.2 Impact of trading costs

We now incorporate trading costs into this framework. Let C_{sp} be the log of cumulative transaction costs incurred over the period. Then the following result holds:

Proposition 2. *If $|\rho|$ is not too large and $b \geq a$, the expected log transaction costs incurred are approximated by:*

$$\begin{aligned} \mathbb{E}[C_{sp}] \;\; \approx \;\; & c\,\mathbb{P}(r_{t-1} \leq a) \; + \; c\,(T-2)\Big[\mathbb{P}(r_{t-1} \leq a,\; r_t \geq b) + \mathbb{P}(r_{t-1} \geq b,\; r_t \leq a) \; + \\ & p_1\mathbb{P}(a < r_{t-1} < b,\; r_t \leq a) + (1-p_1)\mathbb{P}(a < r_{t-1} < b,\; r_t \geq b)\Big], \end{aligned} \tag{9}$$

where p_1 is defined as in Proposition 1. If $b < a$, then the expected transaction costs are approximately:

$$\begin{aligned} \mathbb{E}[C_{sp}] \;\; \approx \;\; & c\,\mathbb{P}(r_{t-1} \leq a) \; + \; c\,(T-2)\Big[\mathbb{P}(r_{t-1} < b,\; r_t \geq b) + \mathbb{P}(r_{t-1} > a,\; r_t \leq a) \; + \\ & p_2\mathbb{P}(b \leq r_{t-1} \leq a,\; r_t \leq a) + (1-p_2)\mathbb{P}(b \leq r_{t-1} \leq a,\; r_t \geq b)\Big], \end{aligned} \tag{10}$$

where p_2 is defined as in Proposition 1.

To validate these results, we estimate the expected log return on stop-loss strategies using simulations for various values of the parameters in the model. We then compare the simulation estimates with the approximations obtained using Propositions 1 and 2; Tables 1.4 and 1.5 in the Appendix report the results. The approximations are very good, with the deviation between the simulated and the approximated values not exceeding 0.8% per year in each case, and not exceeding 0.4% in most cases.

As in the random walk case, we can derive conditions under which the stop-loss strategy beats the buy-and-hold strategy. Suppose we employ a tight stop-loss/start-gain policy and consider the case in which $a=b=0$. Using (4) and (10), we obtain an approximate lower bound on the serial correlation ρ in order to outperform the buy-and-hold strategy:[4]

$$\rho \quad > \quad 1.25\,\frac{\mu + 2c\big(\mathbb{P}(r_{t-1} \le 0, r_t \ge 0) + \mathbb{P}(r_{t-1} \ge 0, r_t \le 0)\big)}{\sigma}\,. \tag{11}$$

It is clear that ρ has to be positive. When $\rho>0$ and $\mu=0$, we have:

$$\mathbb{P}(r_{t-1} \le 0, r_t \ge 0) \ + \ \mathbb{P}(r_{t-1} \ge 0, r_t \le 0) \quad \le \quad \frac{1}{2}\,. \tag{12}$$

As a result, an approximate lower bound for ρ is:

$$\rho \quad > \quad 1.25\,\frac{\mu + c}{\sigma}\,. \tag{13}$$

For U.S. stocks over the 1964–2014 period, the daily mean is, on average, equal to 5.62% of volatility; however trading costs, as a fraction of volatility, are much higher. Assuming transaction costs of 0.2%, for $\sigma = 1\%$ the lower bound on ρ becomes 32.0%, which is very high. It is evident that for a realistic scenario, we need to have not only a high serial correlation, but also a high volatility. For example, for a daily volatility of $\sigma = 4\%$, the lower bound on autocorrelation is 13.3%. This is still high, but serial correlation of this magnitude is not unrealistic, as we will see later in our empirical results.

1.4 Simulation analysis

To develop intuition for our theoretical results, we simulate the performance of the stop-loss strategy for various return-generating processes with various parameters and compare its performance to that of a simple buy-and-hold strategy. The comparison is made in terms of raw returns, certainty equivalent (CE) in a mean-variance framework, skewness, and maximum drawdown. While we vary the return process for the risky asset, we always assume the risk-free asset yields a 0% return.

We consider two different cases for the strategy depending on the horizon used to

[4] We use the fact that $\frac{\sqrt{2\pi}}{2} \approx 1.25$.

measure the cumulative return. In each case, we set the start-gain level at 0% and vary the stop-loss level. The first strategy type is a one-day stop-loss, where $I = J = 1$, and the stop level ranges from -6% to 0%. The second is a two-week stop-loss, so that $I = J = 10$; here the stop level varies from -14% to 0%.

We also incorporate transaction costs into our simulations. We assume a level of 0.2% per trade, which is approximately half of the average spread between the closing bid and ask prices over all stocks in our sample on all days in 2013 and 2014. We use the average over the most recent two years instead of over the full sample period from 1964 to 2014 because trading costs have declined significantly in recent years, and current levels seem more practically relevant than higher historical averages.

We first consider the AR(1) process, and find that high serial correlation and volatility leads to outperformance of the stop-loss strategy. We also consider the regime-switching process of Kaminski and Lo (2014). Our results are consistent with theirs, namely that outperformance is quite rare. For both processes, the two-week stop-loss strategy leads to a more positive skewness and less negative maximum drawdown than the buy-and-hold strategy.

1.4.1 AR(1) process

Recall the specification of an AR(1) process:

$$r_t \;=\; \mu \;+\; \rho(r_{t-1} - \mu) \;+\; \epsilon_t,\; \epsilon_t \sim WN(0, \sigma^2).$$

There are three parameters: the mean μ, the volatility σ, and the serial correlation coefficient ρ. In our simulations, we vary the annualized unconditional volatility from 20% to 50% (an empirically plausible range for individual U.S. stocks) and serial correlation from -20% to 20%. While the serial correlation in daily U.S. returns is close to 0 historically, it can take on more extreme values over short periods. We fix the mean at 10% per year for simplicity. For each unique set of values for these parameters, we run 100,000 simulations over a 252-day horizon. We feel this number of simulations is sufficiently large; standard errors of our estimates are reported in section 1.8.5 in the Appendix.

Figure 1-1 shows the performance statistics for the one-day stop-loss strategy relative to the the buy-and-hold strategy (i.e., we subtract the corresponding statistic for the buy-and-hold strategy from that of the stop-loss strategy). Figure 1-2 displays the comparable results for the two-week strategy.

We see that returns and CE depend positively on serial correlation and volatility, and outperformance occurs only for high values of these two parameters, which is consistent with our model. The magnitude of relative performance is quite dramatic, even for a two-week strategy that uses a longer horizon to make decisions and hence does not trade as frequently. For a -20% serial correlation, the stop-loss strategy can lose up to 30% per year relative to the buy-and-hold strategy.

It is interesting to compare the two types of strategies in terms of skewness and maximum drawdown. The one-day strategy has lower skewness than the buy-and-hold strategy in most cases. One explanation is that using a one-day return to get out of the risky asset dampens the potential upside and hence reduces the right tail of the return distribution, even if there is improvement in the left tail. In contrast, the two-week strategy yields higher skewness in almost all situations because the effect on the right tail is quite marginal, whereas the downside risk is cut, especially if serial correlation is high.

The one-day strategy improves maximum drawdown only in cases of positive serial correlation; the two-week strategy does so for quite a few values of negative correlation as well. This difference is due to the fact that while the one-day strategy cuts downside risk, it also incurs high transaction costs and produces poor returns when serial correlation is negative.

In conclusion, the returns and certainty equivalent of stop-loss strategies depend heavily on serial correlation and volatility, outperforming the buy-and-hold strategy only for high values of these parameters. For an investor with preferences for positive skewness and a lower drawdown, the two-week strategy is quite attractive since it is able to consistently reduce downside risk without incurring too much in trading costs, while maintaining most of the upside potential.

1.4.2 Regime-switching process

We now consider a Markov regime-switching (MRS) process for daily returns:

$$r_t \;=\; \mu_{i_t} \;+\; \sigma_{i_t}\epsilon_t \quad , \quad \epsilon_t \sim WN(0,1), \tag{14}$$

Figure 1-1: One-Day Stop-Loss Strategy Performance for AR(1) Process

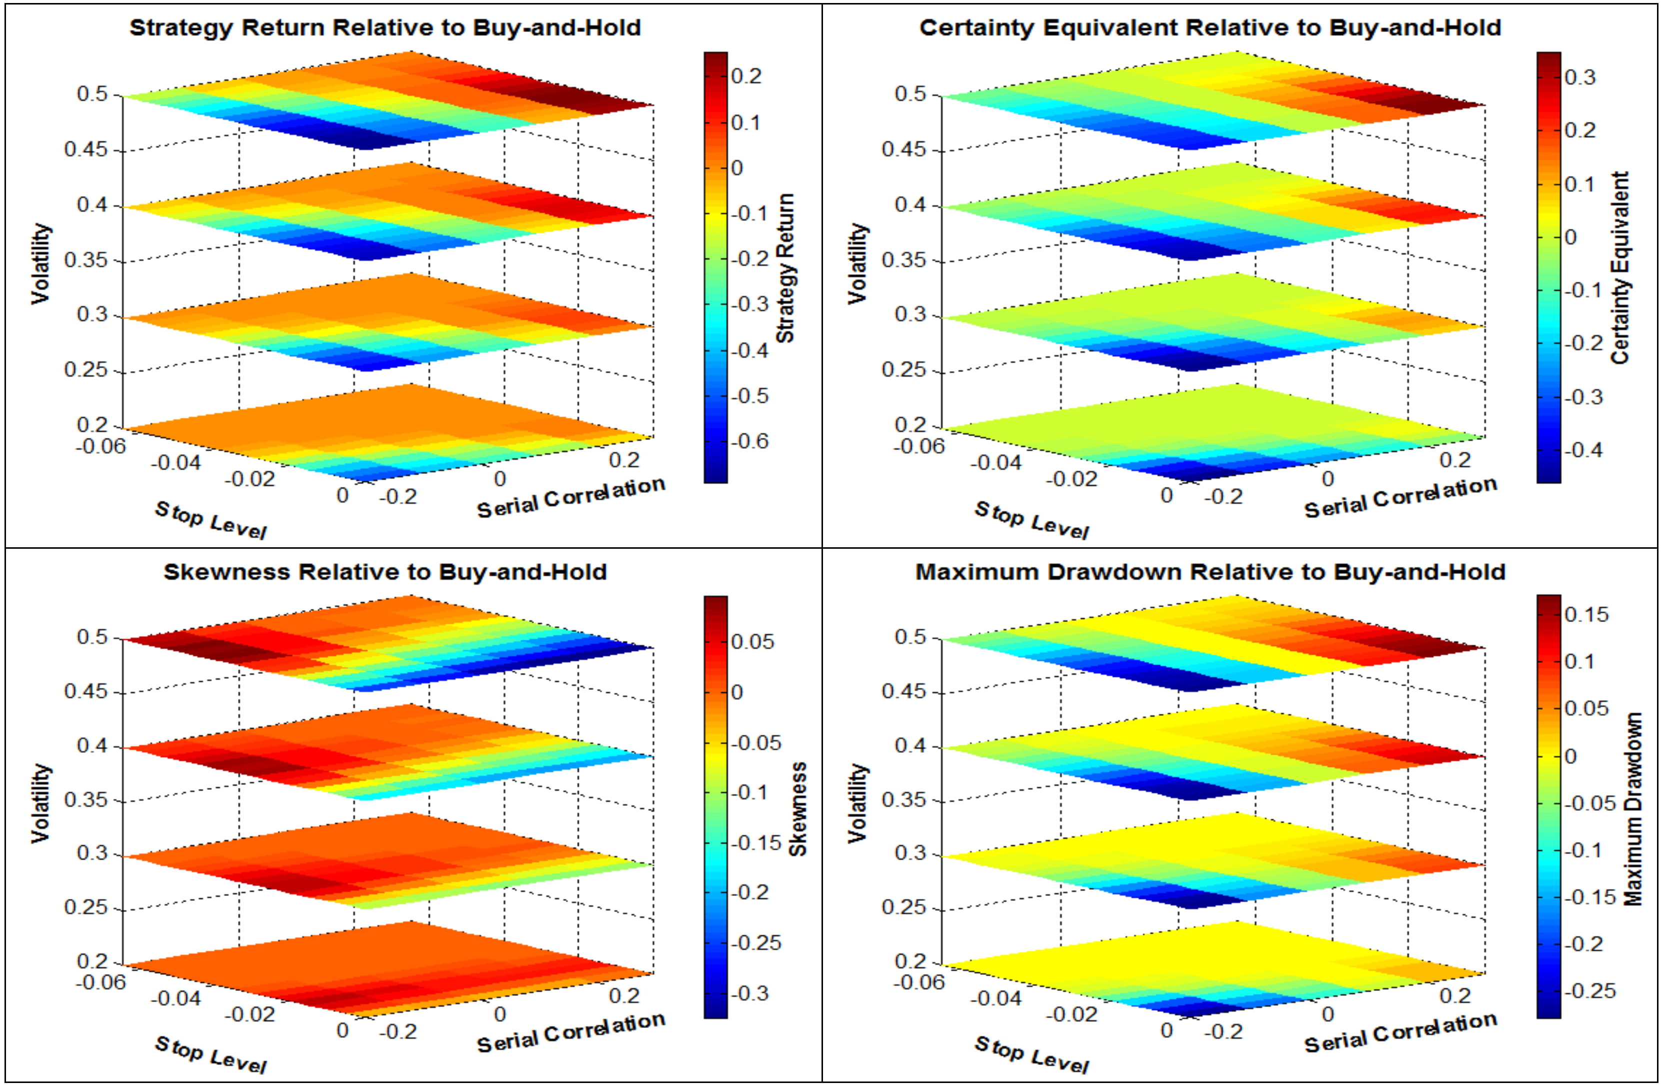

Performance statistics for stop-loss strategy with $I = J = 1$, start level of 0%, and varying stop level. All statistics are measured relative to the buy-and-hold, so that the value for the buy-and-hold strategy is subtracted from that of the stop-loss strategy.

Figure 1-2: Two-Week Stop-Loss Strategy Performance for AR(1) Process

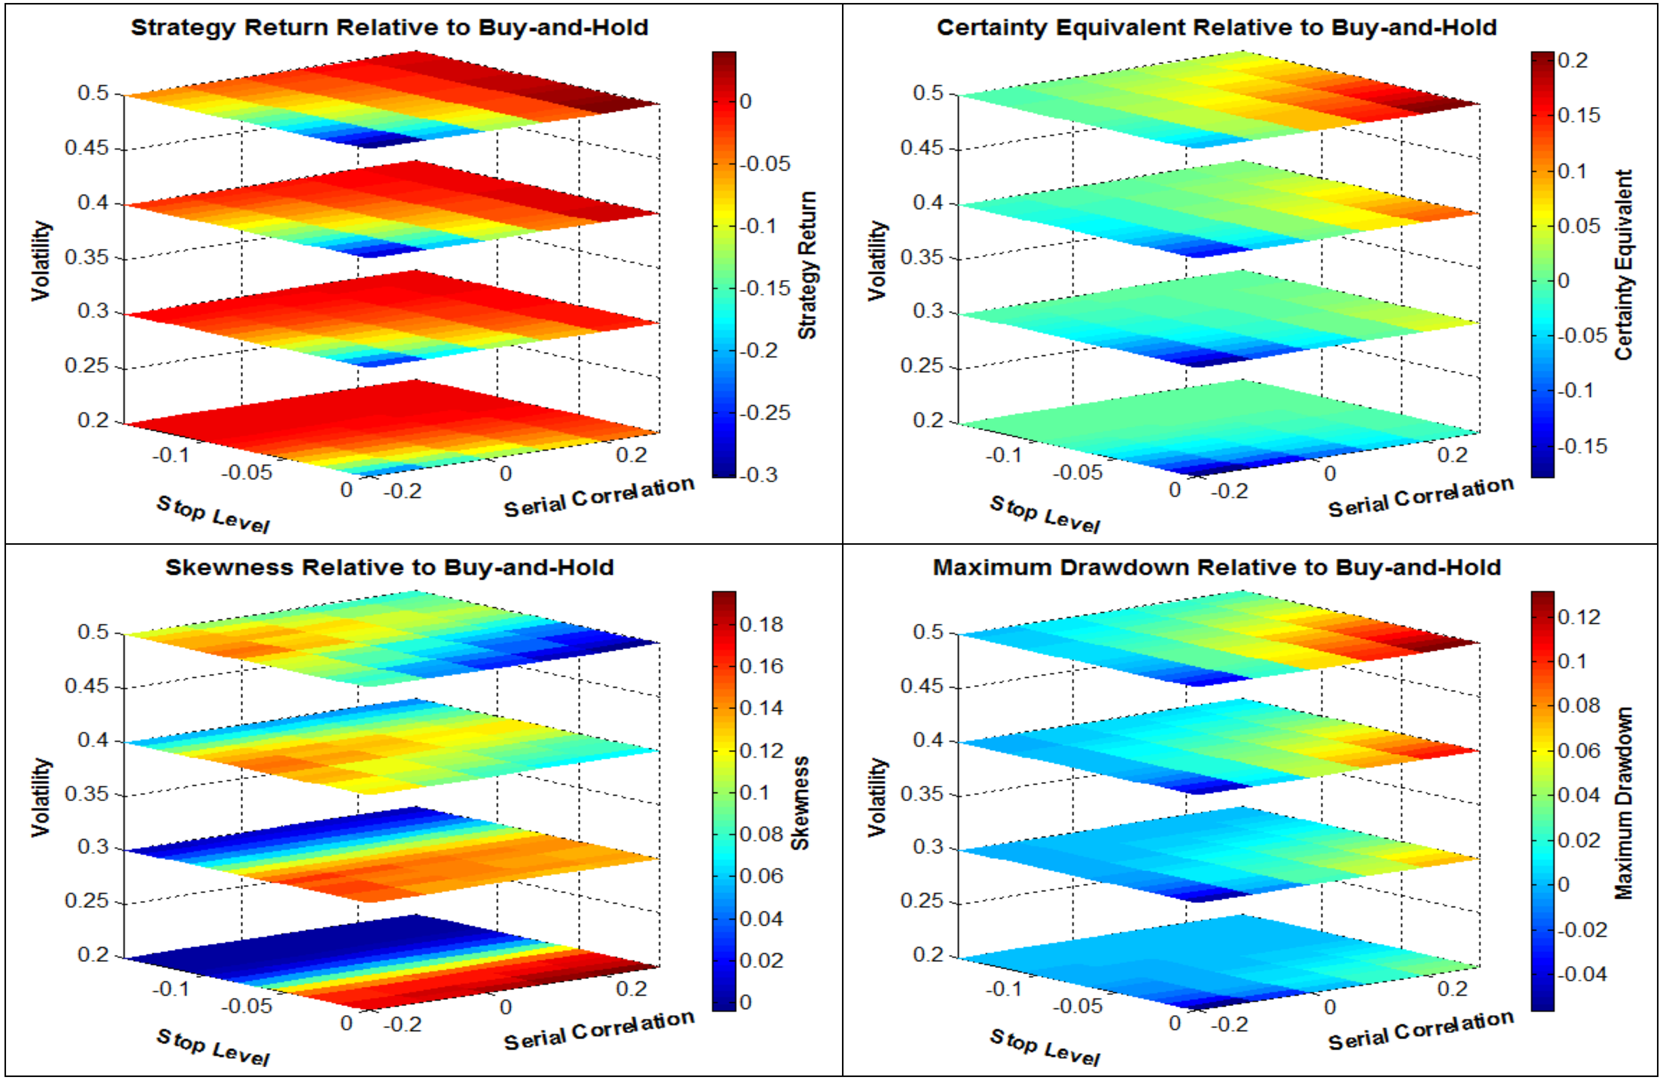

Performance statistics for stop-loss strategy with $I = J = 10$, start level of 0%, and varying stop level. All statistics are measured relative to the buy-and-hold, so that the value for the buy-and-hold strategy is subtracted from that of the stop-loss strategy.

where $i_t \in \{1, 2\}$ is the regime indicator that evolves according to a discrete Markov chain with transition probability matrix P:

$$P = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix}, \tag{15}$$

so that $p_{ij} = \mathbb{P}(i_{t+1} = j | i_t = i)$. In the "bull market" regime (without loss of generality, we assume this is regime 1), the risky asset's returns are distributed as $N(\mu_1, \sigma_1^2)$, and in the "bear market" regime (regime 2), its returns are distributed as $N(\mu_2, \sigma_2^2)$, where $\mu_2 < \mu_1$.

There are three sets of parameters: the means μ_1, μ_2, the volatilities σ_1, σ_2, and the transition probability matrix P. We vary each of them to see how the stop-loss and buy-and-hold strategies perform relative to each other. Table 1.1 lists the parameter values considered, which were chosen to capture a representative range of empirical characteristics for U.S. equities. The more extreme negative values of μ_2 and σ_2 represent stock-market crashes. Both transition probability matrices we consider imply that these negative regimes are not rare but occur much less frequently than the positive regimes (20% and 14% of the time for P_1 and P_2, respectively). For each unique set of parameter values, we run 100,000 simulations over a 252-day horizon; standard errors are provided in the Appendix.

Table 1.1: Parameter Values of MRS Process Simulations

Parameters	Values Considered
(μ_1, μ_2)	(10%, −10%), (15%, −20%), (20%, −30%)
σ_1	20%, 30%
σ_2	40%, 80%
P	$P_1 = \begin{bmatrix} 0.99 & 0.01 \\ 0.04 & 0.96 \end{bmatrix}$, $P_2 = \begin{bmatrix} 0.96 & 0.04 \\ 0.25 & 0.75 \end{bmatrix}$

Table 1.1 lists the parameter values of the simulations for the MRS process. The means μ_1, μ_2 and standard deviations σ_1, σ_2 are annualized. The higher values of σ_2 and the more extreme low values of μ_2 capture stock market crashes. There are two cases for the transition probability matrix P: P_1, when there is little switching out of regimes, and P_2, when there is frequent switching.

Figures 1-3 and 1-4 show the performance statistics for the one-day and two-week stop-loss strategies relative to the buy-and-hold strategy. The stop-loss strategy produces lower returns in almost all cases; it is able to outperform in terms of CE when volatility is high in the bear regime. Performance is better when the expected return in the bear regime is low and when there is little switching between regimes. The intuition behind this is that for the stop-loss strategy to outperform, it must switch out of the risky asset during negative regimes. Further, the negative regime should last for a considerable amount of time so that the relative gain of investing in the risk-free asset will offset transaction costs.

The improvement in CE in situations of high volatility in the bear regime happens for two reasons. First, when volatility is very high, the stop-loss strategy is morely likely to be triggered, correctly divesting when expected returns are negative. Second, the high volatility in the negative regime results in a low CE for a mean-variance investor, so the risk reduction due to holding the risk-free asset produces significant benefits relative to the buy-and-hold strategy.

When it comes to skewness and kurtosis, the results are similar to those of the AR(1) process. The one-day strategy gives lower skewness, while the two-week strategy gives higher skewness in comparison to the buy-and-hold strategy. Furthermore, while the one-day strategy improves maximum drawdown over the buy-and-hold strategy in about half the cases (again, for high volatility in the bear regime), the two-week strategy does so in all situations. We can conclude that it does a good job of managing downside risk.

In summary, the stop-loss strategy beats the buy-and-hold strategy when volatility in the negative regime is high, when the returns in the negative regime are low, and when there is little switching between the two regimes. Using a wider stop results in more frequent outperformance, since stops are then more likely to occur during negative regimes. Finally, the two-week stop-loss strategy offers a consistent reduction in maximum drawdown and a more positive return skewness, whereas this occurs in fewer cases for the one-day strategy.

1.5 Empirical analysis

To gauge the practical relevance of stop-loss strategies, we document their performance when applied to individual U.S. stocks over the 1964–2014 sample period. The raw returns on the strategies are very poor due to excessive transaction costs. When it comes to the certainty equivalent in a mean-variance framework, the results look better due to a reduction in volatility when using a stop-loss strategy in comparison to the buy-and-

Figure 1-3: One-Day Stop-Loss Strategy Performance for MRS Process

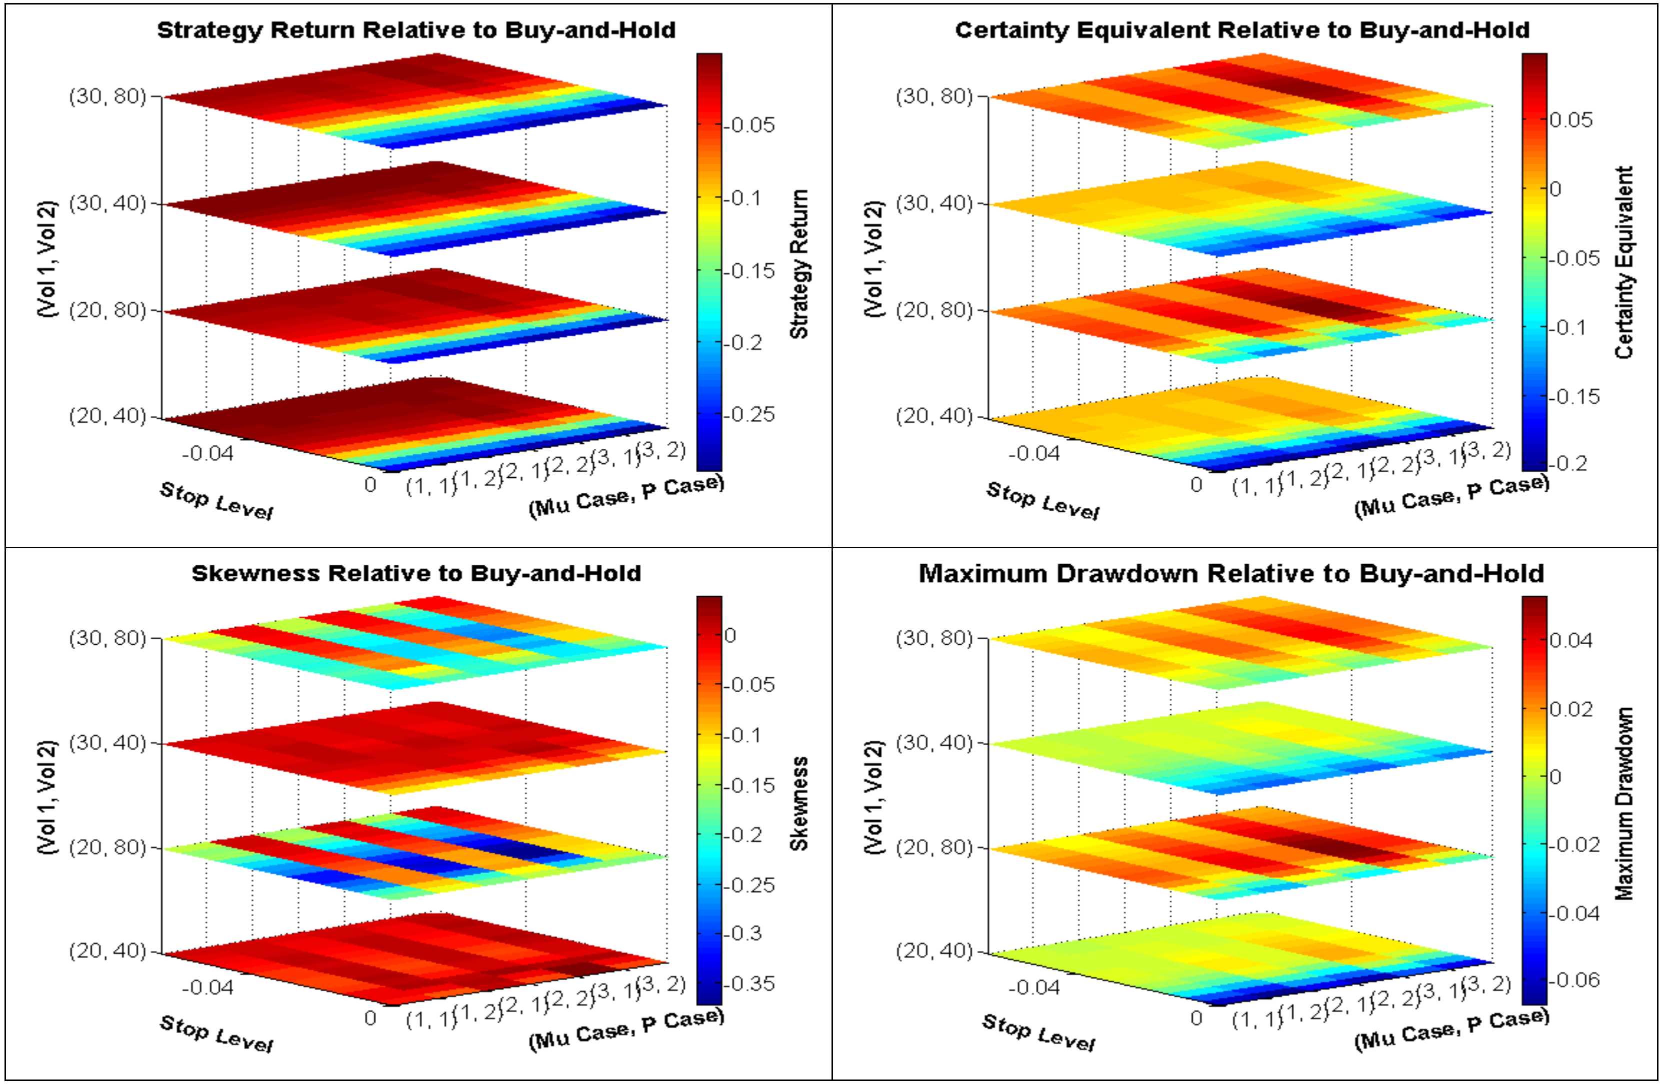

Performance statistics for stop-loss strategy with $I = J = 1$, start level of 0%, and varying stop level. All statistics are measured relative to the buy-and-hold strategy. *Mu Case* corresponds to one of the three possible pairs of means listed in Table 1.1; *P Case* corresponds to one of the two transition matrices listed in the same table.

Figure 1-4: Two-Week Stop-Loss Strategy Performance for MRS Process

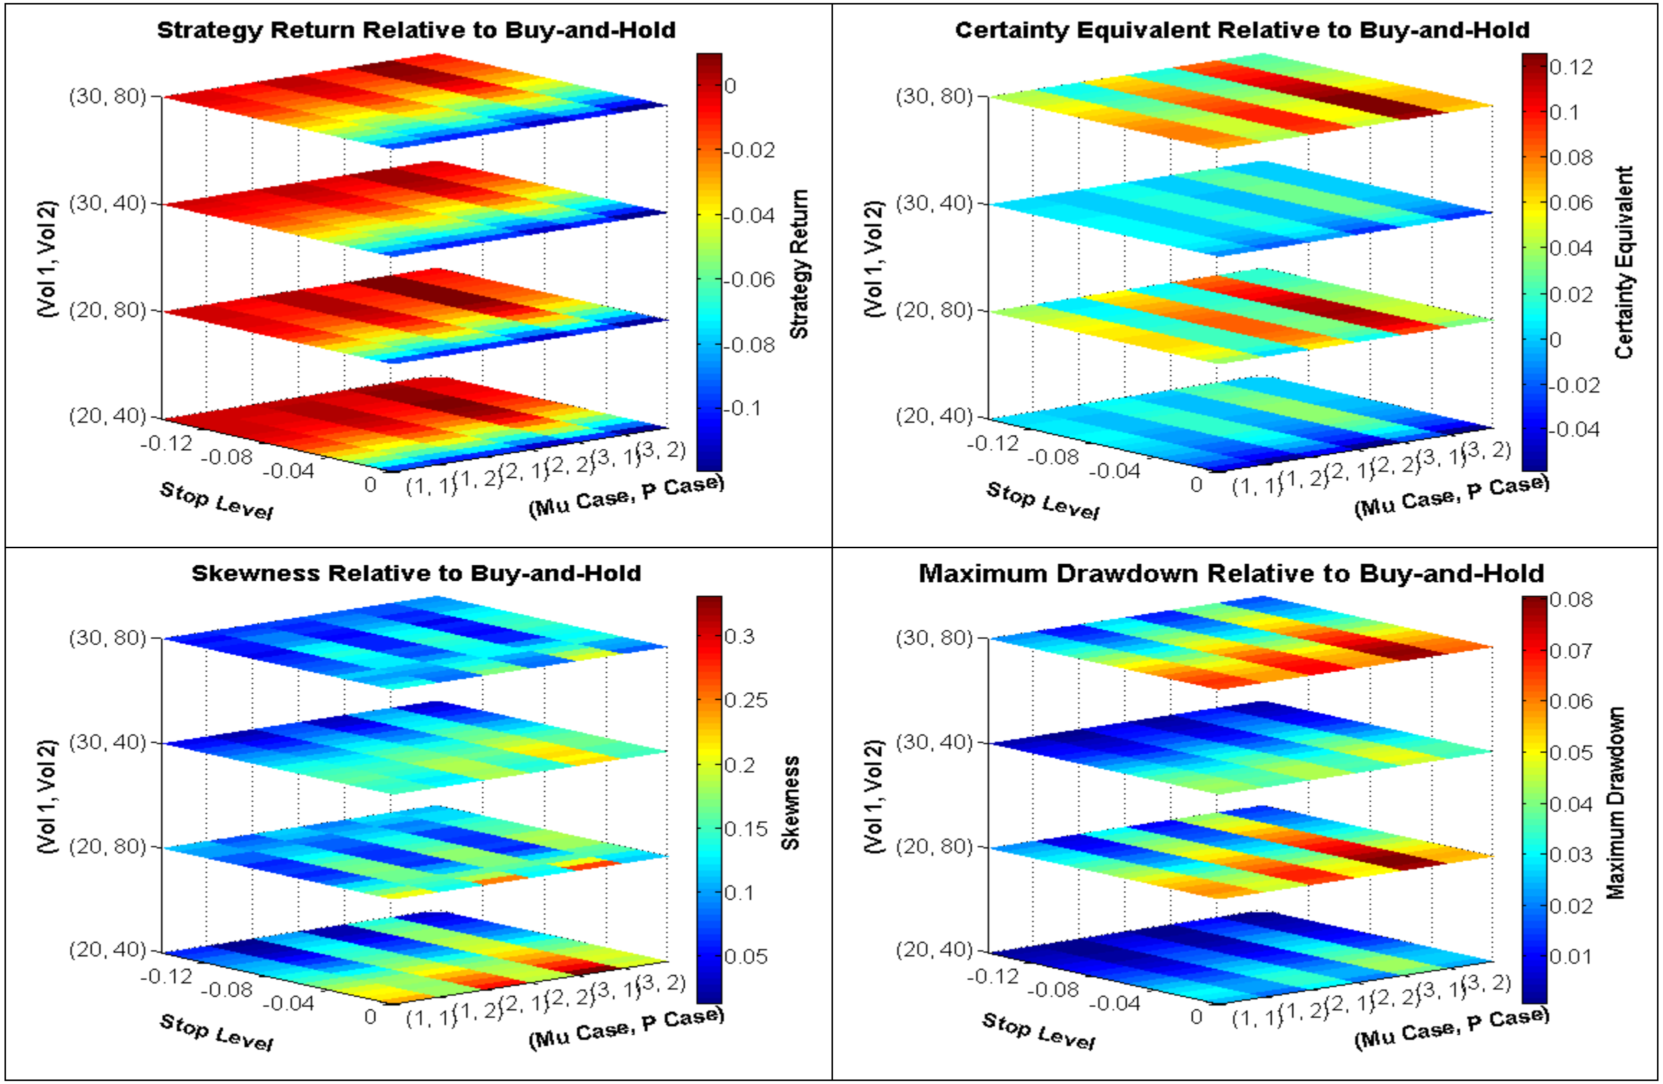

Performance statistics for stop-loss strategy with $I = J = 10$, start level of 0%, and varying stop level. All statistics are measured relative to the buy-and-hold strategy. *Mu Case* corresponds to one of the three possible pairs of means listed in Table 1.1; *P Case* corresponds to one of the two transition matrices listed in the same table.

hold strategy. Finally, we connect the empirical results to our model and demonstrate that historical returns on the strategies exhibit heavy dependence on the interaction term between volatility and serial correlation.

1.5.1 Data and methodology

Stop-loss strategy performance is calculated on a yearly basis. At the start of every year from 1964 to 2014 we use the CRSP database to identify all stocks listed on the NYSE, AMEX, and NASDAQ. We exclude shares of non-U.S. companies, Americus trust components, exchange-traded funds, closed-end funds, and real estate investment trusts. We also remove all stocks with a closing price below $5 on the first trading day of the year, and stocks with less than 250 active trading days before that day. We consider the daily returns of the remaining stocks, adjusted for dividends and stock splits; any missing returns are replaced with zero. Stop-loss strategies are then applied to each of these stocks.

Table 2.1 in the Appendix contains summary statistics for the daily returns in the stock sample. In addition to the statistics for the entire 1964–2014 period, we also provide statistics for consecutive 5-year subperiods (the last subperiod from 2009 to 2014 contains 6 years). As we have posited in Propositions 1 and 2, serial correlation is a very important factor in explaining the performance of tight stop-loss strategies. During the 1964–1988 period, average serial correlation tends to be close to zero or positive. However, from 1989 to 2014, this average is negative over all subperiods. Since average serial correlation is slightly negative over the full period, we expect the historical returns on stop-loss strategies to be inferior to buy-and-hold returns.

We consider two different rates of return on the safe asset. The first is simply 0% all the time; the second is the U.S. 30-day T-bill return. Monthly returns on the T-bill are obtained from Ibbotson Associates and converted to daily returns assuming continuous compounding.

Finally, we incorporate trading costs by assuming that the investor pays one-half of the bid-ask spread whenever a stock is traded, where the spread is calculated with end-of-day closing bid and ask prices from CRSP. However, prior to 1993 these data are missing for a large portion of the stocks. We resort to using estimates from closing prices over the 1964–1992 period from Joel Hasbrouck's website obtained using a Gibbs sampling approach outlined in Hasbrouck (2009).

1.5.2 Strategy performance

We analyze stop-loss strategies with the stop and start horizons ranging from one day to two weeks. We restrict the start-gain level δ to values between 0% and 1.5%, and vary the stop level γ between 0% and -20%. The safe asset is assumed to be the U.S. 30-day T-bill.[5] Figure 1-5 shows strategy performance for the different parameter values.

Tight stop-loss strategies have very poor returns, with the one-day strategy employing a 0% stop losing over -20% per year for all start levels considered, in comparison to the $+15.2\%$ annual gains for the buy-and-hold strategy. This is not surprising since tight stop-loss strategies require a great deal of trading. Strategies with a wider stop-loss limit provide better performance but still underperform the buy-and-hold strategy in terms of raw returns.

Stop-loss strategies do better when we use the certainty equivalent as the basis of comparison. While most strategies still underperform due to worse returns and higher transaction costs, most strategies employing a wide stop limit do as well as the buy-and-hold or even a little better. In particular, two-week strategies with a start gain level between 0.5% and 1.5%, and stop level between 12% and 16% tend to do best. For these strategies, the CE value is between -14.9% and -16.3%, in comparison to -18.3% for the buy-and-hold. The favorable results stem from the volatility reduction upon employing a stop-loss strategy.

For a mean-variance investor, the buy-and-hold strategy is not necessarily optimal. An investor would tend to allocate only a portion of their wealth to the risky asset, based on the perceived mean and variance of the asset, as well as the individual risk aversion. At the same time, the buy-and-hold strategy is widely used in practice and therefore serves as a natural benchmark when looking at stop-loss strategy performance.

Finally, strategies with wider stops have lower skewness, usually worse than the buy-and-hold, whereas strategies with tight stops have higher skewness and outperform. This is because upside is reduced in all cases. However, performance for tight stop-loss strategies is very poor, leading to a significant shift in expected returns and a seemingly shorter left tail. There are quite a few cases when stocks suffer significant intra-day losses, and

[5] In the Appendix, we include results when assuming a 0% risk-free rate of return.

Figure 1-5: Historical Stop-Loss Strategy Performance, 1964-2014; U.S. T-bill as Safe Asset

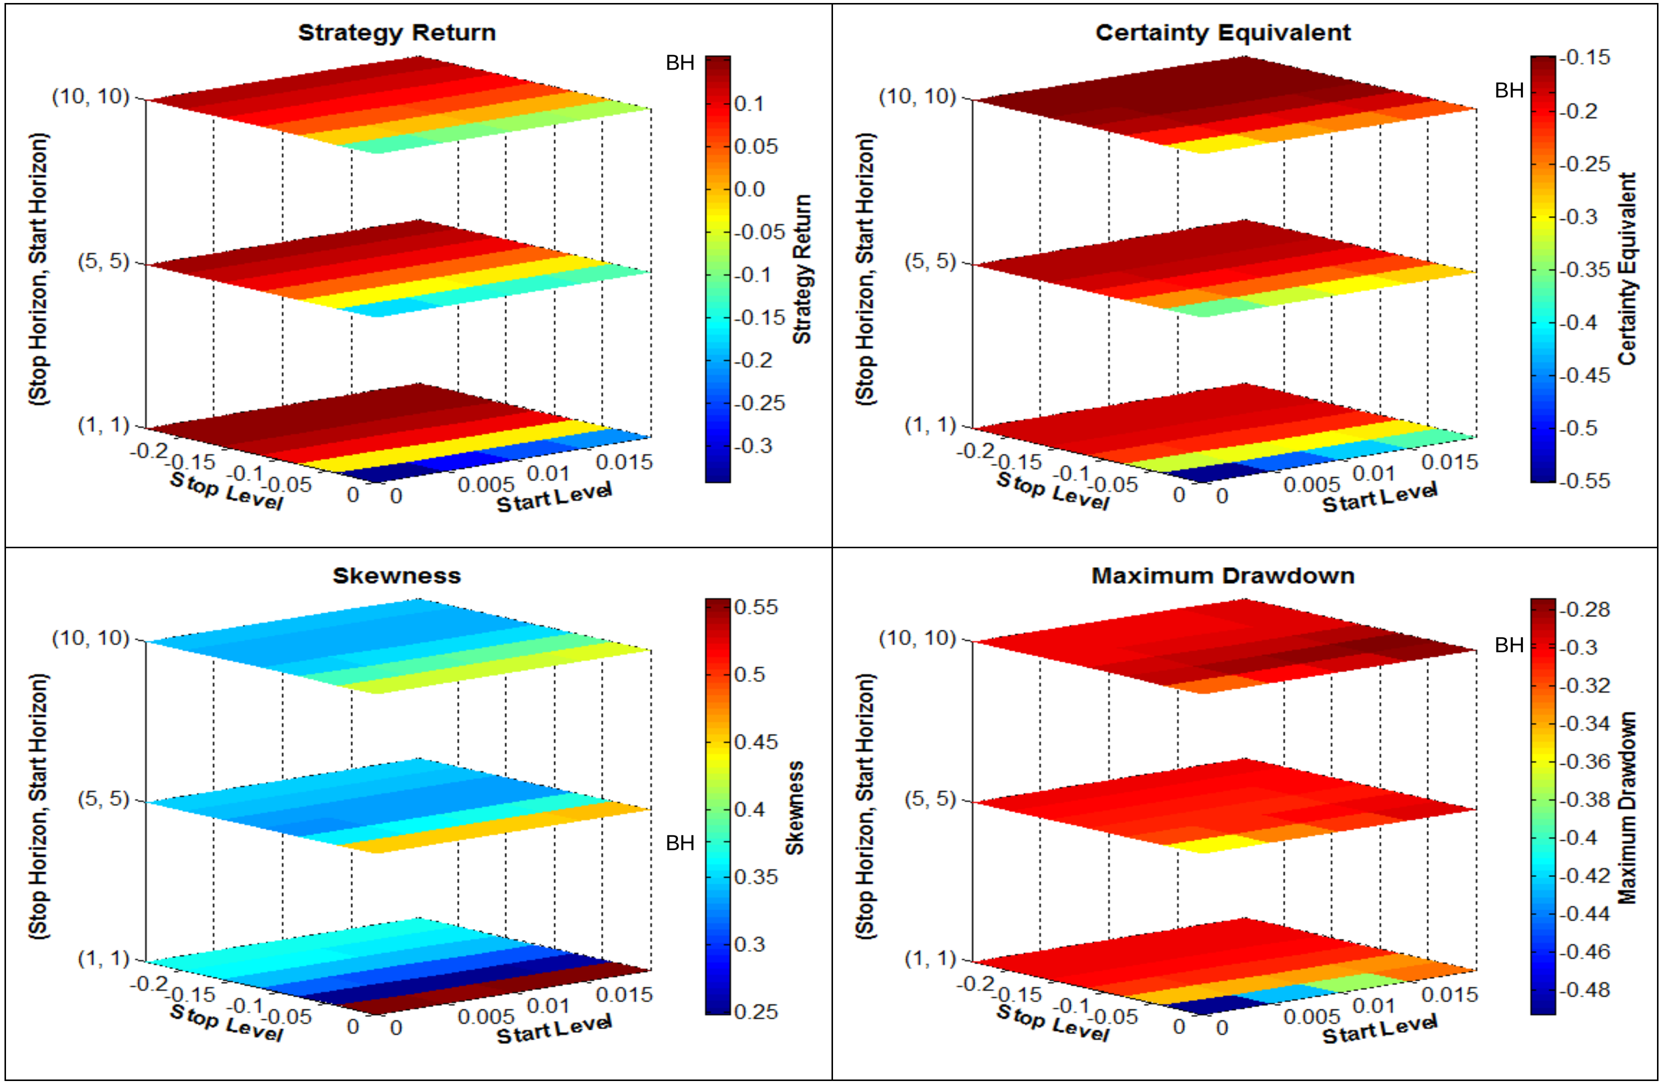

Historical performance of stop-loss strategies with varying levels of stop and start horizons and levels. The “safe" asset is the U.S. 30-day T-bill. All statistics are measured using annual returns and averaged out over all years in the sample. The perfromance statistic for the buy-and-hold is indicated on the colorbar with “BH".

during those cases it is common to observe very high bid-ask spreads (as large as 50% of stock price). This, of course, leads to very poor returns, and in fact quite a few that are worse than −90% over the year, implying a very fat left tail. When it comes to drawdowns, we are able to get improvement over the buy-and-hold strategy, especially when a two-week strategy is used, because the downside is reduced without incurring too much trading cost in the process.

Since many proprietary trading strategies employ tight stops, from now on we focus on strategies with small values for the stop and start levels. We take a "typical" tight stop-loss strategy $\mathcal{S}(-2\%, 0\%, 1, 1)$ (using the U.S. T-bill as the safe asset) and regress its returns in excess of the buy-and-hold strategy on the statistical properties of stock log returns. We control for the time effect as follows. Each "observation" corresponds to the return on the stop-loss strategy applied to a particular stock in a particular year. We employ indicator variables for each of the 51 years in the sample, excluding the last year to avoid collinearity.

We also control for the firm size effect. On the first day of each of the years in the sample, we compute the market capitalization for each of the available stocks, and split the resulting values into deciles. This way we assign a decile to each stock in each year, and we use indicator variables for each of the ten deciles, again excluding the last one.

To summarize, the regression is:

$$R^S_{i,t} - R^{BH}_{i,t} = \alpha + \beta x_{i,t} + \text{controls} + \epsilon, \tag{16}$$

where $R^S_{i,t}$ is the return on the stop-loss strategy for stock i in year t and R^{BH} is the return on the buy-and-hold strategy; $x_{i,t}$ are various functions of stock returns; controls are the time and size controls defined earlier; and ϵ is a random error term.

Table 1.2 contains the results of the regression. The mean (along with time and size controls) is able to explain a significant portion of variation in returns, around 41.7%. Adding volatility and the interaction between volatility and mean increases R^2 by about 6.0%. Adding the remaining regressors boosts R^2 by another 14.1%, a significant increase.

We next consider a set of regressors motivated by our model for the performance of stop-loss strategies when log-returns follow an autoregressive process. Following Propositions 1 and 2, we only include the mean μ and the interaction term $\rho(1) \times \sigma$ between the AR(1) coefficient and volatility as the explanatory variables. Table 1.2 shows that

Table 1.2: Regressions of the Stop-loss Strategy $\mathcal{S}(-2\%, 0\%, 1, 1)$ Relative Return on Statistical Properties of Stock Log Returns

Explanatory Variable	Reg 1	Reg 2	Reg 3	Reg 4	Reg 5	Reg 6	Average
Intercept	15.2%	29.9%	26.6%	25.1%	12.1%	24.6%	N/A
	(20.6)	(42.3)	(38.0)	(41.4)	(18.5)	(40.5)	
Return	-0.388	-0.448	-0.462	-0.510	-0.429	-0.508	3.16%
	(-193.7)	(-127.8)	(-130.3)	(-165.8)	(-240.0)	(-164.9)	
Volatility		-0.660	-0.647	-0.607		-0.591	42.17%
		(-134.9)	(-133.5)	(-143.6)		(-132.7)	
Ret. $\times$ Volatility		-0.072	-0.052	-0.011		-0.013	-3.93%
		(-20.0)	(-14.5)	(-3.5)		(-4.2)	
Skewness / 100			0.008	-0.033		-0.001	0.30%
			(0.1)	(-0.7)		(0.0)	
Kurtosis / 100			0.282	0.192		0.189	11.80%
			(60.4)	(47.1)		(46.5)	
$\rho(1)$				1.218		1.128	-2.95%
				(229.5)		(120.6)	
$\rho(2)$				0.299		0.311	-0.55%
				(32.3)		(33.4)	
$\rho(1)$ $\times$ Volatility					2.246	0.209	-1.87%
					(208.6)	(11.8)	
Adj R-sq	41.70%	47.76%	49.10%	61.84%	54.03%	61.87%	N/A

In table 1.2 we regress the return of the stop-loss strategy $\mathcal{S}(-2\%, 0\%, 1, 1)$ relative to the buy-and-hold strategy on statistical properties of stock log returns. In each regression, we control for time effect and firm size effect by using indicator variables for each year and each market cap decile for the stock at the start of the year. Note that the average value for the log returns is very different than for simple returns in the summary statistics table 2.1 because we are using log returns instead of simple returns.

these two terms explain more than 50% of the variation in returns. Furthermore, the estimated regression coefficients are in line with our model. The coefficient for μ of -0.43 implies highly negative dependence on the expected return of the underlying process, since higher returns hurt relative performance. The coefficient for the interaction term $\rho(1) \times \sigma$ of 2.2 implies significant dependence on serial correlation and volatility. A more detailed discussion of why these coefficients are in line with our model is provided in the Appendix.

If we include all eight summary statistics as regressors, the R^2 only improves marginally, suggesting that our model is able to explain the relative performance of tight stop-loss strategies very well, and that this performance has a close to linear dependence on the product of serial correlation and volatility.

To further investigate the dependence of strategy returns on autocorrelation, in each year we divide the sample of stocks into seven groups based on their realized serial correlation for the year. In each group, we compute the return on the $\mathcal{S}(-2\%, 0\%, 1, 1)$ strategy in excess of the buy-and-hold strategy using the U.S. T-bill as the "safe" asset. Table 1.3 contains the results of this procedure. We find a drastic difference in performance across the realized serial correlation groups. The tight stop-loss strategy applied to stocks with autocorrelation exceeding 15% outperforms the buy-and-hold strategy by 5.5% per year, while it dramatically underperforms by 57% per year for stocks with autocorrelation less than -10%. The overall patterns in this table show that higher autocorrelation leads to significantly better returns. It should be noted that trading costs and positive equity risk premia can cause a tight stop-loss strategy to underperform the buy-and-hold strategy even for stocks with positive serial correlation. The tight stop policy outperforms buy-and-hold only for the stocks in the highest autocorrelation group, and even then only during some of the sub-periods of the entire 1964–2014 sample.

Table 1.3 also contains the excess returns on the strategy in different autocorrelation groups over time. Throughout all of the subperiods, higher autocorrelation gives much better relative performance. The pattern of severe underperformance for stocks with low serial correlation and outperformance for stocks with high serial correlation holds for most subperiods as well.

To develop greater intuition for positive serial correlation in equity returns, we record the proportion of stocks in our sample that fall in a particular serial correlation group

Table 1.3: Returns on the Stop-loss Strategy $\mathcal{S}(-2\%, 0\%, 1, 1)$ in Excess of the Buy-and-hold Stratified by Serial Correlation

	Serial Correlation	1964 - 1968	1969 - 1973	1974 - 1978	1979 - 1983	1984 - 1988	1964 - 2014
Excess Return	-10% or less	-42.9%	-41.3%	-50.2%	-49.2%	-67.8%	-56.6%
	-10% to -5%	-48.3%	-40.1%	-46.1%	-46.7%	-41.1%	-41.5%
	-5% to -0%	-41.7%	-33.7%	-38.8%	-37.2%	-30.4%	-33.7%
	0% to 5%	-33.6%	-23.7%	-30.9%	-29.4%	-21.9%	-25.3%
	5% to 10%	-24.1%	-13.1%	-20.4%	-18.0%	-12.2%	-15.7%
	10% to 15%	-14.9%	-2.1%	-13.5%	-10.1%	-4.6%	-6.1%
	15% or more	-1.6%	13.4%	0.6%	3.2%	5.8%	5.5%
Proportion of Stocks	-10% or less	16.8%	13.1%	11.6%	8.5%	26.6%	27.0%
	-10% to -5%	11.7%	10.1%	9.0%	7.2%	9.1%	11.7%
	-5% to -0%	15.6%	13.3%	13.1%	11.5%	11.9%	14.9%
	0% to 5%	18.0%	17.0%	15.3%	14.3%	13.5%	14.9%
	5% to 10%	15.5%	17.0%	15.5%	15.9%	13.3%	12.3%
	10% to 15%	11.1%	13.3%	14.2%	14.5%	10.6%	8.5%
	15% or more	11.3%	16.2%	21.3%	28.1%	14.9%	10.6%

	Serial Correlation	1989 - 1993	1994 - 1998	1999 - 2003	2004 - 2008	2009 - 2014	1964 - 2014
Excess Return	-10% or less	-87.4%	-83.7%	-63.6%	-36.7%	-45.1%	-56.6%
	-10% to -5%	-51.1%	-57.4%	-46.5%	-18.9%	-22.6%	-41.5%
	-5% to -0%	-41.7%	-51.8%	-40.5%	-11.3%	-13.8%	-33.7%
	0% to 5%	-32.8%	-43.7%	-32.7%	-3.7%	-4.4%	-25.3%
	5% to 10%	-21.5%	-35.3%	-23.6%	4.3%	3.2%	-15.7%
	10% to 15%	-11.8%	-28.3%	-8.6%	13.1%	15.1%	-6.1%
	15% or more	1.8%	-20.3%	5.5%	22.3%	20.9%	5.5%
Proportion of Stocks	-10% or less	40.7%	41.0%	34.1%	30.6%	27.7%	27.0%
	-10% to -5%	9.2%	11.4%	14.0%	16.7%	17.4%	11.7%
	-5% to -0%	12.1%	13.5%	16.1%	19.4%	22.1%	14.9%
	0% to 5%	12.0%	12.9%	15.2%	16.4%	17.8%	14.9%
	5% to 10%	10.9%	10.0%	10.6%	10.0%	9.9%	12.3%
	10% to 15%	7.3%	6.0%	5.7%	4.5%	3.7%	8.5%
	15% or more	7.8%	5.2%	4.4%	2.4%	1.5%	10.6%

Serial correlation is calculated using daily log returns. *Proportion of Stocks* measures the proportion of stocks in each bucket over the particular period. This is done by calculating the proportion of stocks in a particular bucket for each year and then taking the average over all years. We report results from averaging over all years, as well as over all subperiods. Note that there are more buckets for positive values of serial correlation since these are situations when the stop-loss strategy is more likely to outperform the buy-and-hold strategy.

each year and the striking results are given in Table 1.3. During the first half of the 1964–2014 period, many stocks exhibited high autocorrelation: more than 20% of stocks in each subperiod with correlation exceeding 10%, with as much as 43% of such stocks around 1980. However, over the most recent 25 years, there have been much fewer of these stocks: 15.2% in the 1989–1993 period, and less than that in subsequent periods, with only 5.2% of such stocks in 2009–2014.[6]

The pattern is reversed for mean-reverting stocks (i.e., those with low serial correlation). In the 1964–1983 period, there were under 20% of stocks with autocorrelation less than −10%. This proportion jumped to 27% in 1984–1988, then to 41% in 1989–1993, and has stayed above 27% in each of the subsequent subperiods. The explanation for this pattern is beyond the scope of this paper. However, as we have demonstrated, using tight stop-loss strategies provides a simple yet effective way to trade serial correlation; thus there is significant economic value in being able to forecast it.

1.5.3 Delayed stop-loss strategies

As documented in Section 1.5.1, daily U.S. stock returns exhibit slight negative autocorrelation over the 1964–2014 period, with most of it occurring during the past two decades. This suggests that stocks may have a tendency to revert in the short-term following large price movements. This anomaly has been well-documented in the finance literature (e.g., Bremer and Sweeney, 1991; Benou and Richie, 2003; Savor, 2012). As a result, the *delayed* stop-loss strategies $\mathcal{S}^d$ may provide superior returns to their non-delayed counterparts.

Recall that with a delayed stop-loss strategy we wait an extra day to trade out or trade into the risky asset. For example, if the past return over a certain horizon was below a specified threshold on day t, then the strategy would switch to the safe asset at the end of day $t+1$ instead of day t.

Figure 1-6 contains the performance metrics of the delayed stop-loss strategy relative to its non-delayed counterpart using the same specifications. As before, we consider three different past horizon pairs (I, J), stop-loss levels ranging from 0% to −24%, and start-gain levels between 0% and 1.5%. We see that the delayed strategy provides an

[6]We calculate these proportions by adding the percentages of stocks in the buckets for correlation between 10% and 15% and for correlation of 15% or more.

improvement in all cases for the return, CE, and maximum drawdown. The improvement is particularly drastic for the one-day strategy using tight stops, since this is the one for which one-day reversals would be most relevant. For example, one-day strategies using a stop of 0% and start-gain under 1% experience an improvement of 2% to 5% per year when using the delayed specification.

For strategies with wider stops performance gets better when using delays, but it is very marginal. Thus overall, their performance relative to the buy-and-hold would not change drastically, and it would still look similar to Figure 1-5. Finally, we note that delayed strategies usually have better skewness than the non-delayed ones – because upside is more preserved by capturing the positive returns on days following price declines. The only exception is very tight one-day strategies, where skewness decreases; however this is not due to reducing upside, but due to the significant shift in the average (and very negative) strategy return resulting from high transaction costs.

We conclude that using delayed stop-loss strategies marginally improves performance in most cases, and significantly improves returns for the case of tight stops. They are also easier to execute, since the investor can just submit a market-on-close order at the end of the next day rather than trying to submit one right before the close while tracking returns in real-time as with the original stop-loss strategies. Therefore, if an investor does decide to employ stop-loss strategies within our framework, it is generally beneficial to use the delayed specification.

1.6 Optimizing Stop-Loss Level

By now we have a good understanding of how stop-loss strategy performance depends on the underlying market dynamics and strategy parameters. We briefly discuss how to tackle the question of choosing an optimal stop-loss level. We apply the optimization framework to the case of an AR(1) process and investigate how the optimized stop level depends on the process volatility and serial correlation, as well as the investment horizon.

As before, there is a risky asset with log-returns r_t in period t, and a risk-free asset with a constant log-return r_f in each period. An investor trades these two assets over time, with a portion w_t of his portfolio held in the risky asset. There is a fixed cost c (as a percent of notional traded) associated with trading the risky asset.

Figure 1-6: Delayed Stop-Loss Strategy Performance, 1964-2014

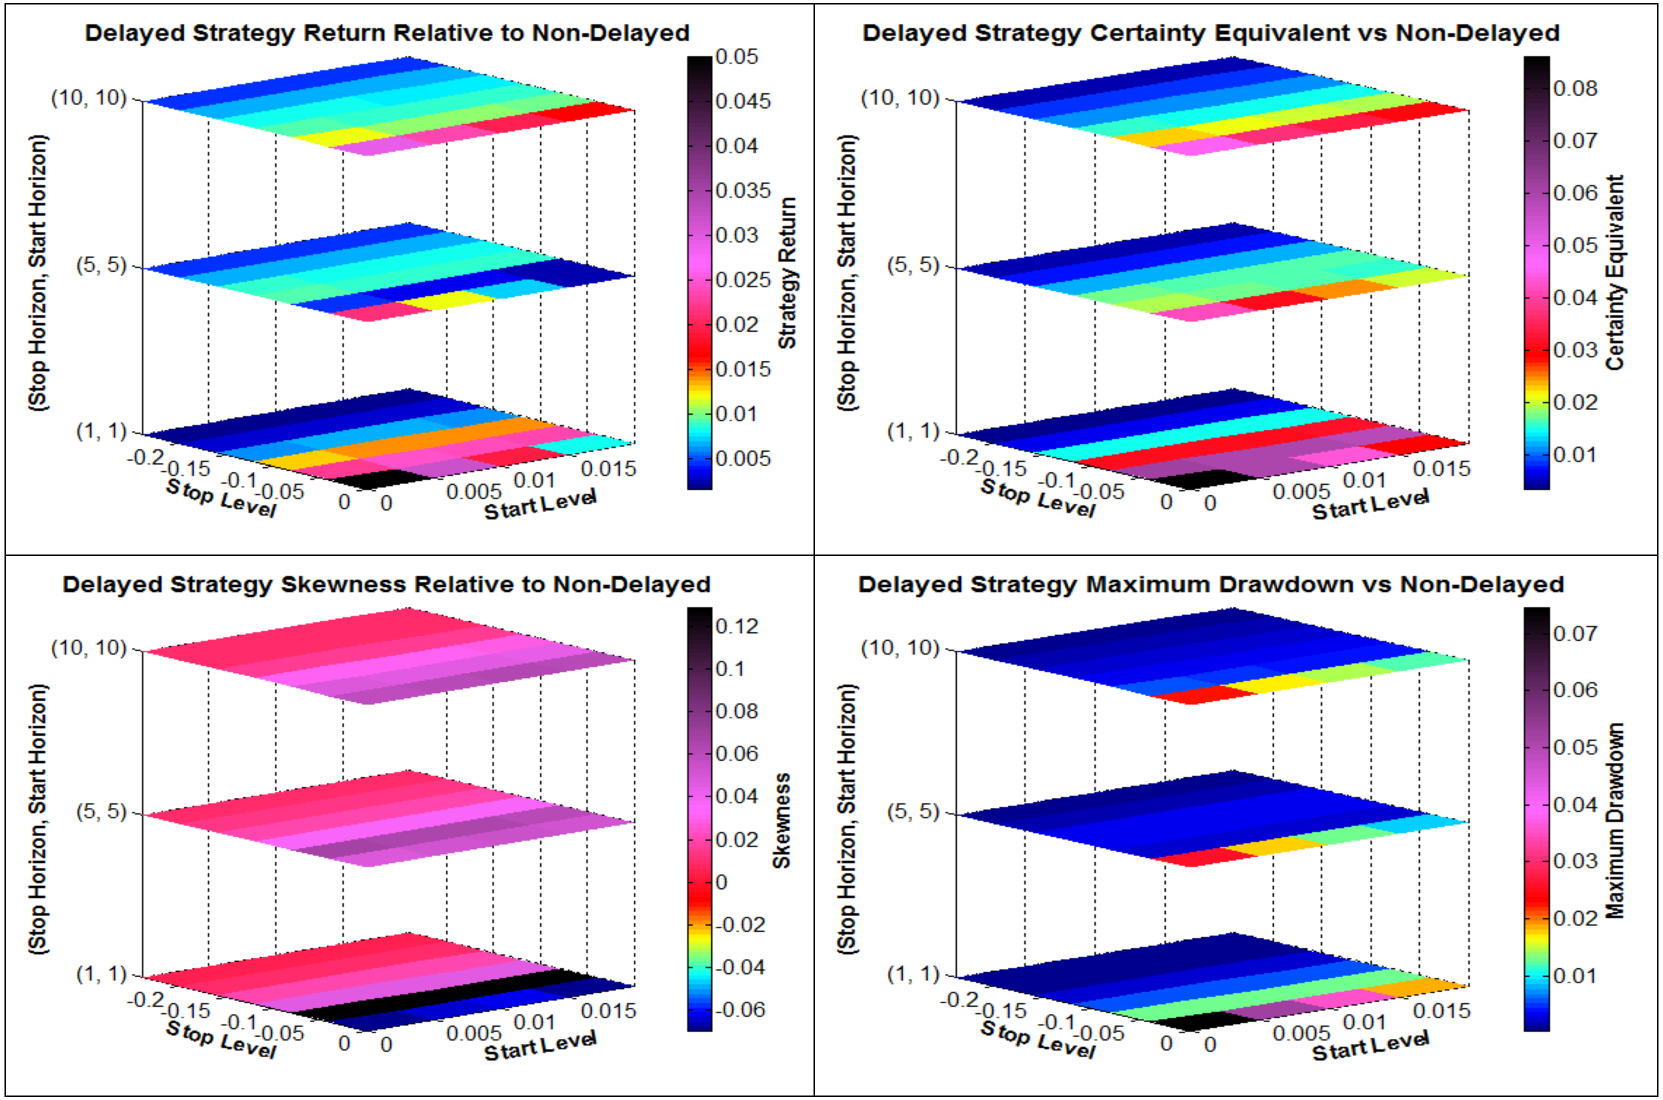

Historical performance of delayed stop-loss strategies relative to the non-delayed ones using the same specifications. Relative performance is measured by taking differences in the corresponding performance metrics. The “safe" asset is the U.S. 30-day T-bill. All statistics are measured using annual returns and averaged out over all years in the sample.

There are T periods, and therefore $T+1$ times $0, 1, \ldots, T$. An investor starts with initial wealth W_0 and can trade at times $t = 0, 1, \ldots, T-1$ by deciding on allocation w_{t+1} during period $t+1$. His objective is to maximize the expected utility of terminal wealth $\mathbb{E}(U(W_T))$.

At time t the investor uses all of the information available up to that point to make his decision on portfolio allocation w_{t+1} during the next period. Define by $V(W_t, w_t, \mathcal{I}_t, t)$ to be the *value function* which is the expected utility of terminal wealth, *conditional* on period t wealth W_t, allocation w_t, and information $\mathcal{I}_t$ up to time t. Accounting for the wealth dynamics, we can write, for $t = 0, 1, \ldots, T-1$:

$$V(W_t, w_t, \mathcal{I}_t, t) = \max_{w \in \mathcal{S}} \mathbb{E}_t \big[V(W_{t+1}, w, \mathcal{I}_{t+1}, t+1) \big] \; s.t. \tag{17}$$

$$W_{t+1} = w \times W_t(1 + r_{t+1}) + \Big[(1-w) \times W_t - c \times |w - w_t| \times W_t \Big] (1 + r_f)$$

where $\mathcal{S}$ is the set of possible values that the allocation can take. After the last period there is no uncertainty, so that $V(W_T, w_T, \mathcal{I}_T, T) = U(W_T)$. The optimal allocation w_{t+1} is defined as:

$$w^*_{t+1} = \underset{w \in \mathcal{S}}{\operatorname{argmax}} \; \mathbb{E}_t \big[V(W_{t+1}, w, \mathcal{I}_{t+1}, t+1) \big] \tag{18}$$

subject to the wealth evolution constraint. There have been several popular approaches in the literature for solving the above problem. One “standard" method is Stochastic Dynamic Programming (SDP), outlined, for example, in Infanger (2006). The idea is to calculate the optimal decision function $V(W_t, w_t, \mathcal{I}_t, t)$ recursively for $t = T-1, \ldots, 0$ at specified grid points for W_t, w_t, and $\mathcal{I}_t$. At each stage t, and each grid point $(W_t, w_t, \mathcal{I}_t)$ the procedure is as follows. We approximate the conditional expectation function:

$$f(w) = \mathbb{E}_t \big[V(W_{t+1}, w, \mathcal{I}_{t+1}, t+1) \big] \tag{19}$$

by drawing a large sample of next period returns (conditional on current information $\mathcal{I}_t$) and averaging the associated values for $V(W_{t+1}, w, \mathcal{I}_{t+1}, t+1)$. These values can be computed, because we know the value function at the next period gridpoints, and we can interpolate in between.

Once we can compute $f(w)$, it suffices to carry out optimization over the set $\mathcal{S}$ to find an approximation for w^*_{t+1} and the corresponding value $f(w^*_{t+1})$ of the value function, which would be an approximation for $V(W_t, w_t, \mathcal{I}_t)$. After dealing with all the gridpoints at time t, we move on to time $t-1$ and repeat.

Brandt et. al (2005) also present an SDP approach, however they propose a novel way to approximate the optimal solution w^*_{t+1}. They perform Taylor expansion on the value function $V(W_t, w_t, \mathcal{I}_t, t)$ and derive the corresponding First Order Conditions. These produce an expression for the optimal solution as an explicit function of partial derivatives of the value function in the next period. The advantage of this approach is that it can deal with more complicated dynamics for the evolution of information $\mathcal{I}_t$ than other simulation-based methods.

Finally, Moallemi and Sağlam (2015) demonstrate how to solve the above problem when considering only the space of *linear dynamic policies*. That is, the next period allocation w_t is restricted to be a linear function of the previous allocation w_{t-1} and *factors* f_s that make up the information set $\mathcal{I}_t$. The authors demonstrate that in that case they can optimize over *all* periods $t = 0, 1, \ldots, T-1$ at once in a single convex optimization problem. The method is shown to be tractable and provides near optimal results.

1.6.1 Applying SDP to Stop-Loss Problem

We now go into more detail about how to solve the optimization problem (17) in the context of our framework. As before, suppose there are two assets: the risky asset with a return r_t during period t and the risk-free asset yielding a constant return r_f each period. The risky asset log returns $\tilde{r}_t$ follow an AR(1) process:

$$\tilde{r}_{t+1} \;=\; \mu \;+\; \rho(\tilde{r}_t - \mu) \;+ \sigma\epsilon_t,\ \epsilon_t \sim WN(0,1) \tag{20}$$

where μ is the mean, σ is the volatility, and ρ is the serial correlation coefficient. We assume the investor knows this is indeed the right model and also its parameters μ, σ, and ρ. At the end of each period $t = 0, 1, \ldots, T-1$ he is able to observe the return r_t for that period (as well as the previous periods); he chooses his allocation w_{t+1} for the next period.

The investor chooses his allocation using the forecast distribution $r_{t+1}|\mathcal{I}_t$ conditional on the information $\mathcal{I}_t$ available as of the end of period t. Because returns follow an AR(1) process and the investor knows this, then it suffices for him to just consider the return in the previous period and look at the corresponding distribution $r_{t+1}|r_t$. This significantly reduces the dimensionality of the problem.

We assume the investor has initial wealth W_0 and has a quadratic utility function $U(W) = W - \lambda W^2$. The allocations in each period are restricted to be either 0% or 100%.

We are now ready to solve (17) using approximate SDP. We first create a grid of nodes (W_t, w_t, r_t) for each time t. For allocation w_t we consider just the values 0% and 100% (since these are all the possible allocations the investor can use).We consider 100 evenly spaced values for log returns $\tilde{r}_t$ in the range $[\mu - k \times \widehat{\sigma}, \mu + k \times \widehat{\sigma}]$, where k is a scaling factor and $\widehat{\sigma} = \sigma/\sqrt{1-\rho^2}$ is the unconditional volatility of the AR(1) process. With $k = 3$, about 99.7% of the log return values sampled under the unconditional distribution will fall into the interval we use, which is very good coverage. From the log returns we easily obtain simple returns r_t as $r_t = \exp(\tilde{r}_t) - 1$.

For wealth W_t we consider 101 equally spaced values in the range $[W_0 \exp(\mu t - k\sigma\sqrt{t}, W_0 \exp(\mu t + k\sigma\sqrt{t})]$, where k is again the scaling factor. Note that this range increases with t to accommodate for the fact that as time passes, the distribution of values of wealth becomes more and more dispersed. The motivation for the functional form of the bounds comes from the evolution of Geometric Brownian Motion after time t, assuming drift μ and volatility σ. Of course, here things are a bit more complicated due to the serial correlation ρ in returns; however for a small value of ρ (under 20% in absolute value) and a large value of k (we use $k = 3$) we will again be covering a large proportion of the sample of possible values of wealth. Note that Brandt et al. (2006) use the same functional form.

Within this set of values W_t we insert another value of W_0, in order to compare the optimal stop-loss level as a function of investment horizon, while keeping initial wealth constant. Thus in the end we use 102 values of W_t in the nodes at each time t.

Once the nodes are well-defined, we outline how to do optimization at each node. We solve the optimization problem recursively at all nodes at time $T-1$, then at all nodes at time $T-2$, and so on. Take an arbitrary node (W_t, w_t, r_t). There are two cases.

Case 1: $t = T - 1$. In this case we know the exact formula for the next period value function:

$$V(W_{t+1}, w_{t+1}, r_{t+1}) = V(W_T, w_T, r_T) = U(W_T) = W_T - \lambda W_T^2 \tag{21}$$

As discussed in the previous section, we need to calculate the conditional expectation function $f(w)$ in (19), where we condition on past period return r_t. To do this, we resort to a “direct" density approximation, where we take 100 evenly spaced values $\tilde{r}_{t+1,s}$ in the range $[\mu_c - k \times \sigma_c,\ \mu_c + k \times \sigma_c]$, where $k = 3$ is a scaling factor giving very good coverage of the distribution, and μ_c, σ_c are the parameters for the conditional distribution of $\tilde{r}_{t+1}$:

$$\tilde{r}_{t+1}|\tilde{r}_t \sim N(\mu_c, \sigma_c^2);\ \ \mu_c = \mu + \rho(\tilde{r}_t - \mu),\ \ \sigma_c = \sigma$$

For each of the values $\tilde{r}_{t+1,s}$ we calculate the next period wealth W_{t+1} and the corresponding value function V as in (21). These are then combined using a weighted average, where the weights correspond to the densities for $\tilde{r}_{t+1,s}$ under the distribution $N(\mu_c, \sigma_c^2)$, and scaled to add to 100%. This gives us $f(w)$.

The final step is optimization. This is easy, because we know $w \in \{0\%, 100\%\}$, so it suffices to just calculate $f(0\%)$ and $f(100\%)$ and take the larger value.

Case 2: $t \leq T - 2$. In this case we know the value function $V(W_{t+1}, w_{t+1}, r_{t+1})$ at specific nodes in the next period. We carry out the exact same procedure as in Case 1 by calculating next period wealth W_{t+1} for each value $\tilde{r}_{t+1,s}$. The only new part is that the next period triplet $(W_{t+1}, w_{t+1}, r_{t+1})$ may not hit an exact node in the next period. If that happens, we perform interpolation between the nodes (we resort to linear interpolation). Note that we need to only interpolate over (W_{t+1}, r_{t+1}) because $w_{t+1} \in \{0\%, 100\%\}$ gives us full coverage of all the possible values for w. After performing the interpolation we can calculate $f(w)$ and again optimize.

1.6.2 Optimization Results

We apply the above optimization approach to an investment problem with a horizon of $T = 20$ periods. We assume each period is a month[7].The annual mean is assumed to be 10%, and we consider different cases for annualized volatility, ranging from 10% to 50%. We also consider different cases for the serial correlation ρ, in the range $[-20\%, 20\%]$. These parameters are the same as in the simulations for an AR(1) process we carried out

[7]We considered daily frequency as well, and results were similar.

earlier.

We look at two cases for transactions costs: $c = \{0.2\%, 1.0\%\}$, as well as two cases for the risk aversion coefficient: $\lambda \in \{0.05, 0.1\}$. The transactions costs levels are motivated by our empirical analysis from before, where a level of 1.0% is closer to the historical average, whereas a level of 0.2% is closer to the average cost of trading at the close over the 2013–2014 period. The risk aversion coefficient choice is rather arbitrary and just done to be small enough so that there is no satiation, where if the past return is very high, the investor decides to cut allocation due to the nature of his utility function. We assume $W_0 = 1$.

While we solve the optimization at each node (W_t, r_t, w_t), we only focus on the nodes where $W_t = W_0 = 1$ so that we have a more fair comparison in behavior across time. We further restrict to nodes where $w_t = 100\%$, since this gives us the case when the investor was fully invested in the previous period. Fixing these two variables, the optimal allocation in the next period depends just on the past return r_t:

$$w^*_{t+1}(W_t, w_t, r_t) = g(r_t)$$

The aim is to show, using the dependence of $g(r_t)$ on r_t, that stop-loss behavior arises at optimality. Recall from before, that for a particular time t, the possible values of r_t considered for the nodes are independent of t. Let us label these values from smallest to largest as $\widehat{r}_1, \widehat{r}_2, \ldots, \widehat{r}_K$.

If we find $g(\widehat{r}_k) = 100\%$ for $k = 1, 2, \ldots, K$, this means that it is always optimal to stay invested, and no stop-loss level is needed. If we find $g(\widehat{r}_k) = 0\%$ for $k \leq L$ and $g(\widehat{r}_k) = 100\%$ for $k > L$, then this means it is optimal to get out of the risky asset for all values of past return below a certain threshold. This is exactly the case when a stop-loss strategy arises, and we can define $\widehat{r}_L$ as the stop-loss level.

Of course, other cases for g are possible, and in those cases a stop-loss strategy is no longer optimal.

Our results are as follows. For $\rho < 0$, stop-loss behavior does not arise, and instead it is the case that $g(\widehat{r}_k) = 0$ for $k \geq L$ and $g(\widehat{r}_k) = 1$ for $k < L$ – i.e. for a large enough return, it is optimal to get out of the risky asset. This makes sense, since if there is negative serial correlation, then a large (positive) past period return is expected to be followed by a relatively significant negative return the next period. Because we are interested in stop-loss behavior, we do not go into more detail with these results.

For $\rho \geq 0$ we find two cases. Either it is optimal to always stay invested, or we do indeed find that stop-loss behavior is optimal. Figure 1-7 plots the stop-loss level as a function of the underlying model parameters. We see that it is optimal to not use a stop-loss when serial correlation is low and/or volatility is low. It makes sense intuitively that if serial correlation is low (and positive), then even if the return in the previous period is very negative, it will have only a small effect on the risky asset return the next period. Because the unconditional expected return is positive, it makes sense to stay invested whenever the conditional expected return is close to it, as is the case when serial correlation is small.

The explanation for why low volatility results in not using a stop-loss is as follows. Since the expected return is fixed, then the return in the next period has to be very negative *in magnitude* for an investor to get out of the risky asset. But if volatility is low, the probability this very negative return occurs is very small – so even for those situations using a stop-loss is better, they occur so rarely (i.e. less than 0.3% of the time) that it's almost like not using a stop-loss at all.

We also see that with higher serial correlation and higher volatility, it becomes optimal to employ a stop-loss strategy; furthermore, the optimal stop-loss level increases and gets closer to 0%. It is important to note that this level has a heavier dependence on serial correlation than on volatility. We also find that as the time horizon increases, the optimal stop-loss level is lower (keeping serial correlation and higher volatility constant), and in some cases it becomes optimal to not use a stop-loss at all. The explanation behind this is due to the fact that using a stop-loss reduces the upside of the terminal wealth, which is particularly important when the investment horizon is long.

Finally, we investigate the dependence of the stop-loss level on the investor risk aversion and transactions costs level. From comparing the two charts on the left of Figure 1-7 and on the right of the figure, we see that the stop level is wider when transactions costs are higher. This makes sense since with higher transactions costs we should not be trading as much and so should set a wider stop-loss level. From comparing the two charts at the top of Figure 1-7 and at the bottom of the figure, we see that the stop level is wider for lower risk aversion – again, this makes sense, since the investor cares about the upside more and is ok with suffering larger losses in wealth.

We conclude that employing a stop-loss strategy is indeed optimal when considering an “all-or nothing" class of dynamic investment strategies. The optimal stop-loss level is increasing in the serial correlation and the volatility and decreasing in the investment horizon. Higher transactions costs and lower risk aversion both result in a lower stop-loss level.

1.7 Conclusion

From both analytical and empirical perspectives, stop-loss strategies can improve investment performance in certain circumstances. Our theoretical results show that the log return on a tight stop-loss strategy is close to linear in the interaction term between return volatility and autocorrelation when returns follow an AR(1) process. We further validate this result by solving for the optimal stop-loss level in a dynamic setting and show that this level is increasing in volatility and serial correlation of the underlying returns process. When returns follow a regime-switching process, wider stop-loss policies will outperform buy-and-hold only in the case when volatility is low in the bull regime and high in the bear regime. And, of course, transaction costs have a significant impact on performance, especially when a tight stop-loss level is used.

When applied to individual U.S. stock returns from 1964 to 2014, we find that stop-loss strategies employing a tight stop produce significantly lower returns in comparison to buy-and-hold strategies. Most of this poor performance stems from very high transaction costs. At the same time, the strategies offer a large reduction in volatility, which in a few cases leads to outperformance over the buy-and-hold strategy in terms of certainty equivalence for a mean-variance investor.

We also explore how stop-loss strategies affect the skewness and maximum drawdown of returns. In our simulations, most two-week strategies are able to provide an improvement over buy-and-hold strategies because they successfully reduce downside risk while preserving upside potential. However, the empirical results using historical stock returns are not as good due to high transaction costs. Nevertheless, some strategies still perform on par with buy-and-hold strategies in terms of skewness and drawdowns.

A closer investigation of the historical performance of tight stop-loss strategies shows that the mean and the interaction term between volatility and serial correlation explain over 50% of the variation in returns on these strategies in excess of the buy-and-hold strategy, and that the estimated regression coefficients are consistent with our analytical results. The trading costs on these strategies are very high, and as a result, a high serial correlation (typically over 10%) is necessary to outperform the buy-and-hold strategy. We also find a striking pattern that over the first half of the 1964–2014 period, stocks tend to exhibit positive serial correlation, while over the second half this correlation turned significantly negative. The mean reversion in returns over the past two decades helps

Figure 1-7: Optimal Stop-Loss Level for AR(1) Process

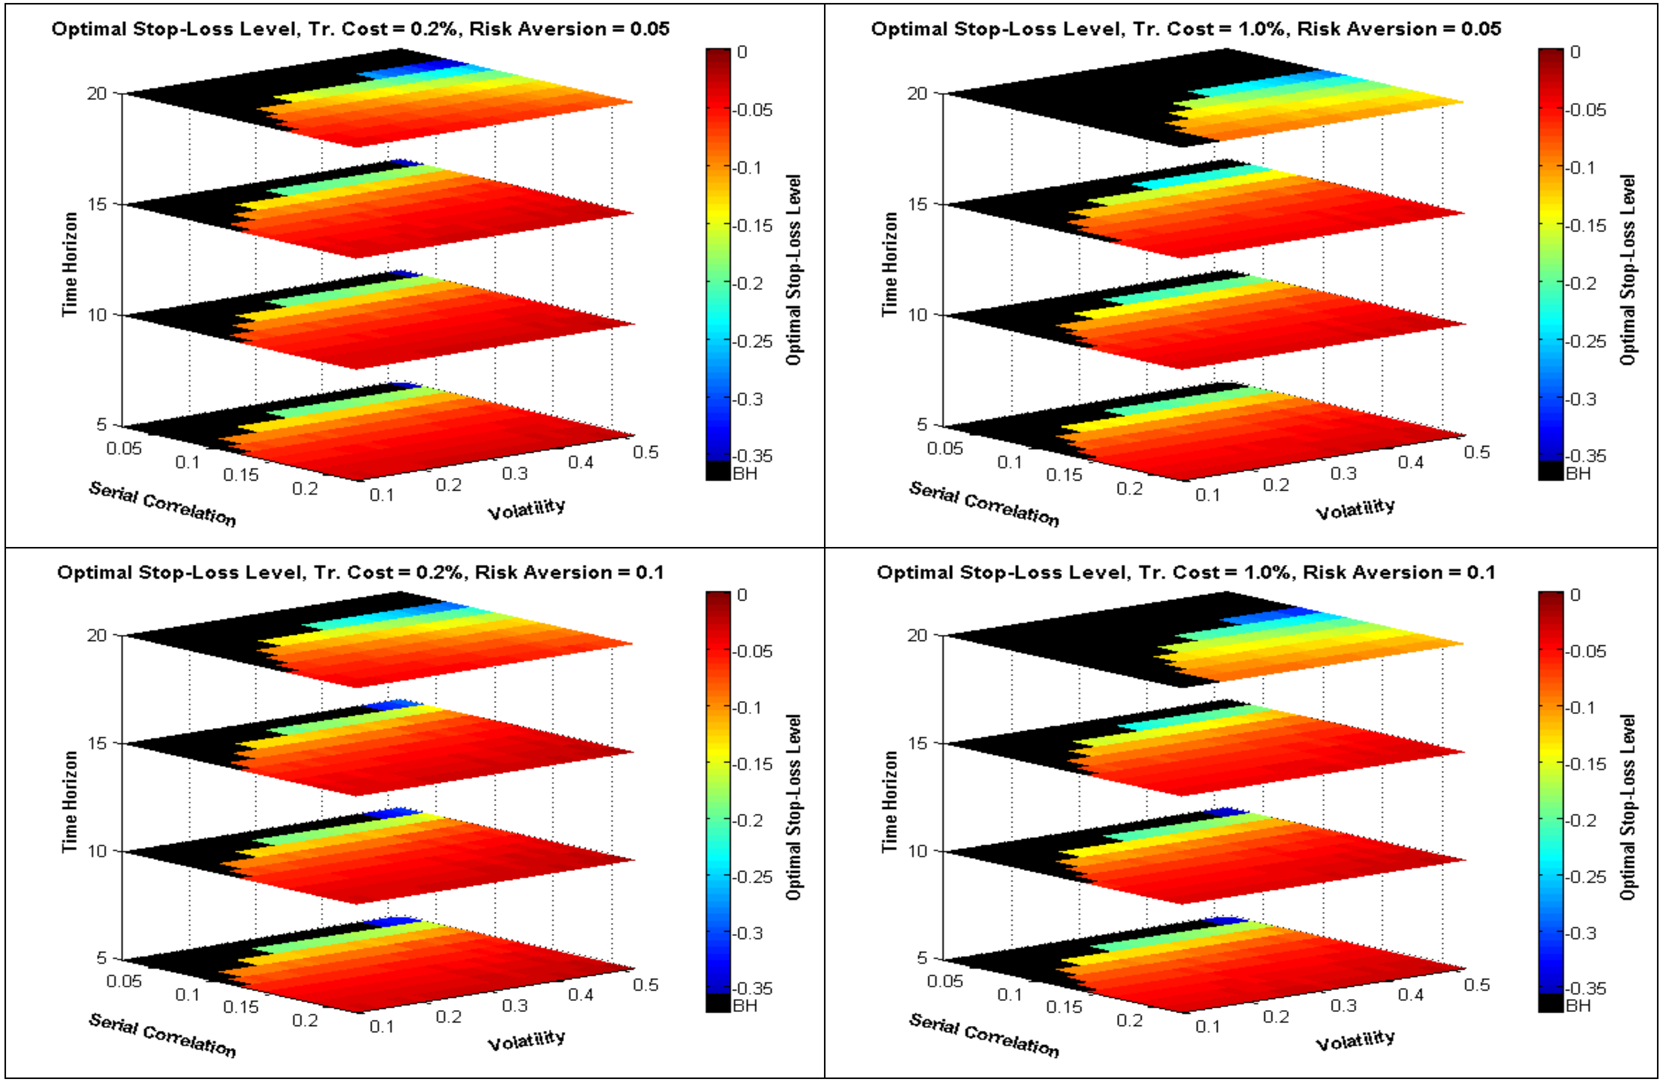

Optimal stop-loss level for a dynamic strategy allocating between a risky and a risk-free asset. The risky asset log returns are assumed to follow an AR(1) process, so that the annual mean of returns is 10%, the annual volatility varies from 10% and 50%, and the serial correlation varies from 0% to 20%. We consider two values for the transactions costs: 0.2% and 1.0%, and two values for the risk aversion coefficient in the quadratic utility function: 0.05 and 0.1. We consider monthly frequency for the returns, and consider different values of the investment horizon, ranging from 5 to 20 periods. The “BH" black area corresponds to a case when it is optimal to always stay invested and to not use a stop-loss.

explain why tight stop-loss strategies have done poorly.

These results clarify the role that stop-loss strategies can play in modern portfolio management. While buy-and-hold portfolios are attractive low-cost passive vehicles for long-term investors, their performance can be improved if asset-price dynamics are more complex than the standard random walk model.

1.8 Appendix

In this Appendix, we provide proofs of propositions 1 and 2 in Sections 1.8.1–1.8.2, the definition of a delayed stop-loss policy in Section 1.8.3, and more detailed discussion of the strategy return regression results in Section 1.8.4. We also include several supplemental tables and figures.

1.8.1 Proof of Proposition 1

For an AR(1) process, the unconditional distribution of log returns is:

$$r_t \sim N(\mu, \tilde{\sigma}^2), \text{ where } \tilde{\sigma}^2 = \frac{\sigma^2}{1-\rho^2} \ . \tag{22}$$

If $|\rho|$ is not too large, e.g., $|\rho| \leq 0.3$, then:

$$|\text{corr}(r_t, r_{t+k})| = |\rho^k| \leq 0.09 \text{ for } k \geq 2 \ . \tag{23}$$

We can thus assume that r_t is approximately independent of $r_{t-2}, r_{t-3}, \ldots$

Case 1: $b \geq a$.

We have $s_1 = 1$ and for $t \geq 2$:

$$s_t = \begin{cases} 1 \text{ if } r_{t-1} \geq b \\ 1 \text{ if } a < r_{t-1} < b \text{ and } s_{t-1} = 1 \\ 0 \text{ otherwise.} \end{cases} \tag{24}$$

For t not too small, we have:

$$\mathbb{P}(s_t = 1 | a < r_t < b) \approx \mathbb{P}(s_{t-1} = 1 | a < r_{t-1} < b). \tag{25}$$

Therefore, in view of (25):

$$\begin{aligned}
\mathbb{P}(s_t = 1|a < r_t < b) &= \mathbb{P}(r_{t-1} \geq b|a < r_t < b) \ + \ \mathbb{P}(s_{t-1} = 1|a < r_{t-1} < b, \ a < r_t < b) \times \\
&\quad \ \mathbb{P}(a < r_{t-1} < b|a < r_t < b) \\
&\approx \mathbb{P}(r_{t-1} \geq b|a < r_t < b) \ + \ \mathbb{P}(s_{t-1} = 1|a < r_{t-1} < b) \times \\
&\quad \ \mathbb{P}(a < r_{t-1} < b|a < r_t < b) \\
&\approx \mathbb{P}(r_{t-1} \geq b|a < r_t < b) \ + \ \mathbb{P}(s_t = 1|a < r_t < b) \times \\
&\quad \ \mathbb{P}(a < r_{t-1} < b|a < r_t < b). \qquad (26)
\end{aligned}$$

We can then define $p_1 = \mathbb{P}(s_t = 1|a < r_t < b)$ and derive the expression for it:

$$\begin{aligned}
p_1 &= \mathbb{P}(s_t = 1|a < r_t < b) \approx \frac{\mathbb{P}(r_{t-1} \geq b|a < r_t < b)}{1 - \mathbb{P}(a < r_{t-1} < b|a < r_t < b)} \\
&= \frac{\mathbb{P}(r_{t-1} \geq b|a < r_t < b)}{\mathbb{P}(r_{t-1} \geq b|a < r_t < b) \ + \ \mathbb{P}(r_{t-1} \leq a|a < r_t < b)}. \qquad (27)
\end{aligned}$$

The above expression can be easily approximated by simulating $r_{t-1} \sim N(\mu, \tilde{\sigma}^2)$ and the innovation $\epsilon_t \sim N(0, \sigma^2)$.

We now have:

$$\mathbb{E}(s_t r_t) = \mathbb{P}(r_{t-1} \geq b)\mathbb{E}(r_t|r_{t-1} \geq b) + \ \mathbb{P}(a < r_{t-1} < b)\mathbb{E}(s_t r_t|a < r_{t-1} < b) \ . \qquad (28)$$

We approximate each term individually for t not too small:

$$\begin{aligned}
\mathbb{P}(r_{t-1} \geq b)\mathbb{E}(r_t|r_{t-1} \geq b) &\approx \mathbb{P}(r_{t-1} \geq b) \times \mu \ + \\
&\quad \rho \int_b^{\infty} (x - \mu) \frac{1}{\sqrt{2\pi}\tilde{\sigma}} \exp\left(-\frac{(x-\mu)^2}{2\tilde{\sigma}^2}\right) dx \\
&= \Phi(\frac{\mu - b}{\tilde{\sigma}})\mu \ + \ \rho \frac{\tilde{\sigma}}{\sqrt{2\pi}} \exp\left(\frac{(\mu - b)^2}{2\tilde{\sigma}^2}\right) \qquad (29)
\end{aligned}$$

and:

$$
\begin{aligned}
& \mathbb{P}(a < r_{t-1} < b)\mathbb{E}(s_t r_t | a < r_{t-1} < b) \\
\approx\ & \mathbb{P}(a < r_{t-1} < b)\mathbb{E}(r_t | a < r_{t-1} < b,\ s_{t-1} = 1) \times p_1 \\
=\ & p_1 \mathbb{P}(a < r_{t-1} < b)\mu \ + \ p_1 \int_a^b (x-\mu)\frac{1}{\sqrt{2\pi}\tilde{\sigma}} \exp\left(-\frac{(x-\mu)^2}{2\tilde{\sigma}^2}\right) dx \\
=\ & p_1(\Phi(\frac{b-\mu}{\tilde{\sigma}}) - \Phi(\frac{a-\mu}{\tilde{\sigma}}))\mu \ + \ p_1 \rho \frac{\tilde{\sigma}}{\sqrt{2\pi}}(\exp\left(-\frac{(a-\mu)^2}{2\tilde{\sigma}^2}\right) - \exp\left(-\frac{(b-\mu)^2}{2\tilde{\sigma}^2}\right)),
\end{aligned}
$$

where p_1 is given by (27). Combining the above two expressions, we get the result for the first case.

Case 2: $b < a$.

This case is similar. We have:

$$
s_t = \begin{cases} 1 \text{ if } r_{t-1} > a \\ 1 \text{ if } b \leq r_{t-1} \leq a \text{ and } s_{t-1} = 0 \\ 0 \text{ otherwise.} \end{cases} \tag{30}
$$

Note that (25) still holds for t not too small. It follows that:

$$
\begin{aligned}
\mathbb{P}(s_t = 1 | b \leq r_t \leq a) \ &= \ \mathbb{P}(r_{t-1} > a | b \leq r_t \leq a) \ + \ \mathbb{P}(s_{t-1} = 0 | b \leq r_{t-1} \leq a, b \leq r_t \leq a) \times \\
& \qquad \mathbb{P}(b \leq r_{t-1} \leq a | b \leq r_t \leq a) \\
&\approx \ \mathbb{P}(r_{t-1} > a | b \leq r_t \leq a) \ + \ \mathbb{P}(s_{t-1} = 0 | b \leq r_{t-1} \leq a) \times \\
& \qquad \mathbb{P}(b \leq r_{t-1} \leq a | b \leq r_t \leq a) \\
&\approx \ \mathbb{P}(r_{t-1} > a | b \leq r_t \leq a) \ + \ \mathbb{P}(s_t = 0 | b \leq r_t \leq a) \times \\
& \qquad \mathbb{P}(b \leq r_{t-1} \leq a | b \leq r_t \leq a).
\end{aligned} \tag{31}
$$

Define $p_1 = \mathbb{P}(s_t = 1 | b \leq r_t \leq a)$, $p_2 = \mathbb{P}(s_t = 0 | b \leq r_t \leq a)$, then:

$$
p_1 \ \approx \ \mathbb{P}(r_{t-1} > a | b \leq r_t \leq a) + p_2 \mathbb{P}(b \leq r_{t-1} \leq a | b \leq r_t \leq a). \tag{32}
$$

Similarly we get:

$$p_2 \;\approx\; \mathbb{P}(r_{t-1} < b | b \leq r_t \leq a) + p_1 \mathbb{P}(b \leq r_{t-1} \leq a | b \leq r_t \leq a). \tag{33}$$

Solving the above system of equations, it follows that:

$$p_1 \approx \frac{1-q_2}{2-q_1-q_2}, \; p_2 \approx \frac{1-q_1}{2-q_1-q_2}, \tag{34}$$

where

$$q_1 = \mathbb{P}(r_{t-1} > a | b \leq r_t \leq a), \; q_2 = \mathbb{P}(r_{t-1} < b | b \leq r_t \leq a). \tag{35}$$

The above expression can be easily approximated by simulating $r_{t-1} \sim N(\mu, \tilde{\sigma}^2)$ and the innovation $\epsilon_t \sim N(0, \sigma^2)$.

We now proceed in a similar way to case 1:

$$\mathbb{E}(s_t r_t) \;=\; \mathbb{P}(r_{t-1} > a)\mathbb{E}(r_t | r_{t-1} > a) + \mathbb{P}(b \leq r_{t-1} \leq a)\mathbb{E}(s_t r_t | b \leq r_{t-1} \leq a). \tag{36}$$

We approximate each term individually:

$$\begin{aligned}\mathbb{P}(r_{t-1} > a)\mathbb{E}(r_t | r_{t-1} > a) &\approx \mathbb{P}(r_{t-1} > a) \times \mu \;+ \\ &\quad \rho \int_a^{\infty} (x-\mu) \frac{1}{\sqrt{2\pi}\tilde{\sigma}} \exp\left(-\frac{(x-\mu)^2}{2\tilde{\sigma}^2}\right) dx \\ &= \Phi(\frac{\mu - a}{\tilde{\sigma}})\mu \;+\; \rho \frac{\tilde{\sigma}}{\sqrt{2\pi}} \exp\left(-\frac{(\mu-a)^2}{2\tilde{\sigma}^2}\right)\end{aligned} \tag{37}$$

and:

$$\begin{aligned}& \mathbb{P}(b \leq r_{t-1} \leq a)\mathbb{E}(s_t r_t | b \leq r_{t-1} \leq a) \\ \approx\; & \mathbb{P}(b \leq r_{t-1} \leq a)\mathbb{E}(r_t | b \leq r_{t-1} \leq a, \, s_{t-1} = 0) \times p_2 \\ =\; & p_2 \mathbb{P}(b \leq r_{t-1} \leq a)\mu \;+\; p_2 \int_b^a (x-\mu) \frac{1}{\sqrt{2\pi}\tilde{\sigma}} \exp\left(-\frac{(x-\mu)^2}{2\tilde{\sigma}^2}\right) dx \\ =\; & p_2(\Phi(\frac{a-\mu}{\tilde{\sigma}}) - \Phi(\frac{b-\mu}{\tilde{\sigma}}))\mu \;+\; p_2 \rho \frac{\tilde{\sigma}}{\sqrt{2\pi}} (\exp\left(-\frac{(b-\mu)^2}{2\tilde{\sigma}^2}\right) - \exp\left(-\frac{(a-\mu)^2}{2\tilde{\sigma}^2}\right)),\end{aligned}$$

where p_2 is given by (34). Combining the above two expressions, we get the result for the second case. ■

1.8.2 Proof of Proposition 2

The technique is similar to Proposition 1. Let $k_t = 1$ if a trade is made at the end of day t; $k_t = 0$ otherwise. We have:

$$\begin{aligned}\mathbb{P}(k_t = 1) &= \mathbb{P}(s_{t+1} = 1,\ s_t = 0) \ +\ \mathbb{P}(s_{t+1} = 0,\ s_t = 1) \\ &= \mathbb{P}(s_t = 0 | r_t \geq b)\mathbb{P}(r_t \geq b) \ +\ \mathbb{P}(s_t = 1 | r_t \leq a)\mathbb{P}(r_t \leq a)\ . \end{aligned} \tag{38}$$

Case 1: $b \geq a$.

We have:

$$\mathbb{P}(s_t = 0 | r_t \geq b) \ \approx\ \mathbb{P}(r_{t-1} \leq a | r_t \geq b) \ +\ \mathbb{P}(s_{t-1} = 0 | a < r_{t-1} < b)\mathbb{P}(a < r_{t-1} < b | r_t \geq b)$$

and

$$\mathbb{P}(s_t = 1 | r_t \leq a) \ \approx\ \mathbb{P}(r_{t-1} \geq b | r_t \leq a) + \mathbb{P}(s_{t-1} = 1 | a < r_{t-1} < b)\mathbb{P}(a < r_{t-1} < b | r_t \leq a).$$

Substituting these into (38) we get:

$$\begin{aligned}\mathbb{P}(k_t = 1) \ &\approx\ \mathbb{P}(r_{t-1} \leq a,\ r_t \geq b) \ +\ \mathbb{P}(r_{t-1} \geq b,\ r_t \leq a) \\ &\quad + p_1\mathbb{P}(a < r_{t-1} < b,\ r_t \leq a) \ +\ (1 - p_1)\mathbb{P}(a < r_{t-1} < b,\ r_t \geq b),\end{aligned}$$

where p_1 is defined as in Proposition 1.

Case 2: $b < a$.

We have:

$$\mathbb{P}(s_t = 0 | r_t \geq b) \ \approx\ \mathbb{P}(r_{t-1} < b | r_t \geq b) \ +\ \mathbb{P}(s_{t-1} = 1 | b \leq r_{t-1} \leq a)\mathbb{P}(b \leq r_{t-1} \leq a | r_t \geq b)$$

and

$$\mathbb{P}(s_t = 1 | r_t \leq a) \ \approx\ \mathbb{P}(r_{t-1} > a | r_t \leq a) + \mathbb{P}(s_{t-1} = 0 | b \leq r_{t-1} \leq a)\mathbb{P}(b \leq r_{t-1} \leq a | r_t \leq a).$$

Substituting these into (38) we get:

$$\begin{aligned}\mathbb{P}(k_t = 1) \;\approx\; & \mathbb{P}(r_{t-1} < b,\ r_t \geq b) \;+\; \mathbb{P}(r_{t-1} > a,\ r_t \leq a) \\ & + p_2\mathbb{P}(b \leq r_{t-1} \leq a,\ r_t \leq a) \;+\; (1 - p_2)\mathbb{P}(b \leq r_{t-1} \leq a,\ r_t \geq b),\end{aligned}$$

where p_2 is defined as in Proposition 1.

Multiplying by c and $T - 2$ we get the result. ■

1.8.3 Definition of Delayed Stop-Loss Strategy

Definition 2. *A* **delayed fixed rolling-window strategy** $\mathcal{S}^d(\gamma, \delta, J, I)$ *is a dynamic asset allocation rule* $\{s_t\}$ *between the risky asset* Q *and the safe asset* F*, such that:*

$$s_t = \begin{cases} 0 \text{ if } R_{t-2}(J) \leq \log(1+\gamma) \text{ and } s_{t-1} = 1 \text{ (exit); otherwise:} \\ 1 \text{ if } R_{t-1}(J) > \log(1+\gamma) \text{ and } s_{t-1} = 1 \text{ (stay in); otherwise:} \\ 1 \text{ if } R_{t-1}(J) \leq \log(1+\gamma) \text{ and } s_{t-1} = 1 \text{ (stay in one more day); otherwise:} \\ 1 \text{ if } R_{t-2}(I) \geq \log(1+\delta) \text{ and } s_{t-1} = 0 \text{ (re-enter); otherwise:} \\ 0 \text{ if } R_{t-1}(I) < \log(1+\delta) \text{ and } s_{t-1} = 0 \text{ (stay out); otherwise:} \\ 0 \text{ if } R_{t-1}(I) \geq \log(1+\delta) \text{ and } s_{t-1} = 0 \text{ (stay out one more day).} \end{cases} \tag{39}$$

When using the delayed strategies, we no longer need to impose the assumption that the stock price does not move significantly right before the close. However, we now assume that a market on close order is executed at a price that is equal to, or at least very close to, the closing price for the day.

1.8.4 Discussion of strategy returns regression results

We show that the estimated coefficients in Regression 5 in Table 1.2 are consistent with the approximation formula in Proposition 1. We have $T = 252$, $a = \log(1 - 0.02), b = \log(1 + 0)$, and from Table 1.2, the daily volatility of log returns is on average equal to $\tilde{\sigma} \approx \dfrac{42.17\%}{\sqrt{252}}$, while the daily mean of log returns is on average equal to $\mu \approx \dfrac{3.16\%}{252}$. Using Proposition 1, the log return on the stop-loss strategy in excess of the buy-and-hold strategy is approximately linear in the annual mean with a coefficient of:

$$\frac{1}{T}\left[1 \;+\; (T-1)(\Phi(\frac{\mu-b}{\tilde{\sigma}}) \;+\; p_1(\Phi(\frac{b-\mu}{\tilde{\sigma}}) \;-\; \Phi(\frac{a-\mu}{\tilde{\sigma}}))) \;-\; T\right]. \tag{40}$$

We already know the empirical values of $T, a, b, \tilde{\sigma}$, and μ. For p_1, we do a further approximation by assuming r_{t-1} and r_t are independent (this is a reasonable assumption if serial correlation is small). We have:

$$\begin{aligned} p_1 &= \frac{\mathbb{P}(r_{t-1} \geq b,\ a < r_t < b)}{\mathbb{P}(r_{t-1} \geq b,\ a < r_t < b) + \mathbb{P}(r_{t-1} \leq a,\ a < r_t < b)} \\ &\approx \frac{\mathbb{P}(r_{t-1} \geq b)}{\mathbb{P}(r_{t-1} \geq b) + \mathbb{P}(r_{t-1} \leq a)} \approx \frac{\Phi(\frac{\mu-b}{\tilde{\sigma}})}{\Phi(\frac{\mu-b}{\tilde{\sigma}}) + \Phi(\frac{a-\mu}{\tilde{\sigma}})}. \end{aligned}$$

Substituting the empirical values, we get $p_1 \approx 69\%$. Substituting into (40), we obtain a "theoretical" coefficient of -0.30, which is close to the empirical value of -0.42.

The coefficient for the interaction term between serial correlation and annual volatility is approximately:

$$\frac{1}{\sqrt{T}}\frac{1}{\sqrt{2\pi}}(T-1)\left[\exp\left(-\frac{(\mu-b)^2}{2\tilde{\sigma}^2}\right) \;+\; p_1(\exp\left(-\frac{(a-\mu)^2}{2\tilde{\sigma}^2}\right) \;-\; \exp\left(-\frac{(b-\mu)^2}{2\tilde{\sigma}^2}\right))\right].$$

Substituting the parameter values, the "theoretical" coefficient becomes 5.20, while the empirical coefficient is 2.25. While the coefficients are a bit different, they have a similar order of magnitude.

The intercept term in the regression corresponds to the transaction costs and the component due to the risk-free rate. From Proposition 2, the transaction costs are approximately:

$$\begin{aligned} &c \times \mathbb{P}(r_{t-1} \leq a) \;+\; c(T-2)[\mathbb{P}(r_{t-1} \leq a,\ r_t \geq b) \;+\; \mathbb{P}(r_{t-1} \geq b,\ r_t \leq a) \\ &\qquad + p_1\mathbb{P}(a < r_{t-1} < b,\ r_t \leq a) \;+\; (1-p_1)\mathbb{P}(a < r_{t-1} < b,\ r_t \geq b)]. \end{aligned}$$

These can be approximated further, assuming r_{t-1} and r_t are independent. The approximation becomes:

$$\begin{aligned}\mathbb{E}(C_{sp}) \approx c \times \Phi(\frac{a-\mu}{\tilde{\sigma}}) + c\,(T-2)\,[2\,\Phi(\frac{a-\mu}{\tilde{\sigma}})\,\Phi(\frac{\mu-b}{\tilde{\sigma}}) \\ + p_1\,(\Phi(\frac{b-\mu}{\tilde{\sigma}}) - \Phi(\frac{a-\mu}{\tilde{\sigma}}))\,\Phi(\frac{a-\mu}{\tilde{\sigma}}) \\ + (1 - p_1)\,(\Phi(\frac{b-\mu}{\tilde{\sigma}}) - \Phi(\frac{a-\mu}{\tilde{\sigma}}))\,\Phi(\frac{\mu-b}{\tilde{\sigma}})].\end{aligned}$$

Assuming a low level of transaction costs of 0.2% per trade so that $c = \log(1 - 0.002)$, we get transaction costs of -15.5% according to our model.

The contribution from the risk-free rate is:

$$r_f \times \left[T - (1 + (T-1)(\Phi(\frac{\mu-b}{\tilde{\sigma}}) + p_1(\Phi(\frac{b-\mu}{\tilde{\sigma}}) - \Phi(\frac{a-\mu}{\tilde{\sigma}})))\right] .$$

Since the average daily log return r_f on the U.S. T-bill is around $\frac{4.88\%}{252}$ over the 1964–2014 period, this contribution comes out to 1.5%. Therefore the regression intercept according to the model should be approximately $-15.5\% + 1.5\% = -14.0\%$. This is a negative number, which would be even lower if we assume higher transaction costs. (Over the 1964–2014 period, the average cost per trade is 0.8%.) However, from Table 1.2 we see that this intercept is positive and significant. We can explain this by the fact that we control for time and size effects. In particular, the indicator variables corresponding to lower market capitalization deciles have very large negative coefficients. This is not surprising, since smaller cap stocks have much higher transaction costs, which translate into poor returns on the tight stop-loss strategy considered here. Because of these negative coefficients, the intercept term in the regression ends up being positive instead of negative as per our model.

1.8.5 Standard Errors in Simulations

Recall that we ran 100,000 simulations for each case of parameter values and strategy specifications, both for the AR(1) process and the Markov Regime-Switching (MRS) process. Here we present the level of standard errors for our estimates of strategy returns, maximum drawdown, skewness, and kurtosis. In particular, we argue that these errors are low enough so that we can be comfortable about making conclusions about the dependence of these performance metrics on parameters as observed in Figures 1-1, 1-2, 1-3, 1-4.

For the AR(1) process, the standard error over all strategies (including the buy-and-hold) for the annual return is at most 0.16%, for the CE – at most 0.37%, for skewness is at most 0.0077, and for maximum drawdown is at most 0.05%. We calculate the standard error for the return by taking the standard deviation over all samples (and scaling appropriately); the same is done for calculating the median of skewness and maximum drawdown (we make sure to use the appropriate multiplier of 1.253 for the median estimate). For the CE we combine the standard errors for the mean and the variance.

For the MRS process, the standard error for all strategies for the annual return is at most 0.16%, for the CE is at most 0.33%, for skewness is at most 0.0077, and for maximum drawdown is at most 0.07%.

Comparing these standard errors to the ranges of observed performance metrics in the figures showing simulation results, it is evident that these errors are very small. Thus using 100,000 simulations for our analysis is sufficient.

1.8.6 Supplemental tables and figures

Table 1.4: Approximation of Stop-loss Strategy $\mathcal{S}(\gamma, 0\%, 1, 1)$ Returns

ρ, γ	-2.0%	-1.0%	0.0%	1.0%	2.0%
-20%	-3.7%	-27.9%	-46.2%	-50.4%	-49.4%
	(-0.4%)	(-0.4%)	(0.0%)	(-0.2%)	(-0.3%)
-10%	0.7%	-17.0%	-32.2%	-37.9%	-38.7%
	(-0.1%)	(-0.2%)	(0.0%)	(-0.1%)	(-0.3%)
0%	4.8%	-6.5%	-18.6%	-25.7%	-28.7%
	(-0.3%)	(-0.2%)	(-0.1%)	(0.0%)	(-0.1%)
10%	9.0%	4.0%	-5.0%	-13.7%	-18.8%
	(0.0%)	(0.0%)	(0.1%)	(0.2%)	(0.3%)
20%	13.9%	15.0%	8.9%	-1.2%	-8.8%
	(-0.2%)	(-0.1%)	(0.0%)	(0.2%)	(0.5%)

Approximation for the expected annual log return of a stop-loss strategy $\mathcal{S}(\gamma, 0\%, 1, 1)$ when using a start-gain level of 0% and a stop-loss level of γ, and using the one-day return to decide whether to stay in or out of the risky asset. We assume an annualized return of 10%, annualized volatility of 20%, risk-free rate of 0%, and transaction costs of 0.2%. The serial correlation ρ ranges from -20% to 20%. The investment horizon is 21 days, corresponding to one month. For each case for (ρ, γ), we run 1,000,000 simulations to estimate the expected log return on the strategy. In parentheses we give the deviation of the simulated value from the theoretical value obtained using the approximation formula in Propositions 1 and 2. The standard error for the simulation estimate in each case does not exceed 0.16%.

Table 1.5: Approximation of Stop-loss Strategy $\mathcal{S}(\gamma, -1\%, 1, 1)$ Returns

ρ, γ	-2.0%	-1.0%	0.0%	1.0%	2.0%
-20%	-2.8%	-26.7%	-46.9%	-52.0%	-50.6%
	(0.0%)	(0.0%)	(0.3%)	(0.6%)	(0.7%)
-10%	1.3%	-16.7%	-34.4%	-42.1%	-43.5%
	(0.0%)	(0.0%)	(0.2%)	(0.2%)	(0.1%)
0%	4.9%	-7.2%	-22.4%	-32.5%	-36.7%
	(0.0%)	(0.1%)	(0.2%)	(0.1%)	(0.0%)
10%	8.6%	2.1%	-10.4%	-22.8%	-29.9%
	(0.0%)	(0.0%)	(0.0%)	(0.1%)	(0.3%)
20%	12.8%	11.8%	1.8%	-12.4%	-22.6%
	(0.0%)	(0.0%)	(-0.1%)	(0.4%)	(0.8%)

Approximation for the expected annual log return of a stop-loss strategy $\mathcal{S}(\gamma, -1\%, 1, 1)$ when using a start-gain level of -1% and a stop-loss level of γ, and using the one-day return to decide whether to stay in or out of the risky asset. We assume an annualized return of 10%, annualized volatility of 20%, risk-free rate of 0%, and transaction costs of 0.2%. The serial correlation ρ ranges from -20% to 20%. The investment horizon is 21 days, corresponding to one month. For each case for (ρ, γ), we run 1,000,000 simulations to estimate the expected log return on the strategy. In parentheses we give the deviation of the simulated value from the theoretical value obtained using the approximation formula in Propositions 1 and 2. The standard error for the simulation estimate in each case does not exceed 0.16%.

Table 1.6: Summary Statistics for Stocks over the 1964–2014 Period

Summary Statistic	1964 - 1968	1969 - 1973	1974 - 1978	1979 - 1983	1984 - 1988	1964 - 2014
Return	29.59%	-7.63%	21.15%	28.24%	11.61%	15.21%
Standard Deviation	32.63%	40.83%	35.15%	35.82%	39.46%	41.31%
Skewness	0.57	0.43	0.52	0.58	0.32	0.38
Kurtosis	5.42	5.48	6.87	7.24	7.12	6.28
Max Drawdown	-22.85%	-38.93%	-28.82%	-26.61%	-29.22%	-29.78%
Serial Correlation	1.03%	3.61%	5.15%	7.62%	-2.01%	-2.10%
Sharpe Ratio	0.85	-0.08	0.71	0.89	0.49	0.52
Certainty Equivalent	10.5%	-37.2%	-1.3%	5.0%	-18.1%	-18.3%
Daily abs(mean)/vol	6.40%	5.08%	6.50%	6.59%	5.46%	5.62%
Number of Stocks	1732	1969	2984	3375	3537	3184

Summary Statistic	1989 - 1993	1994 - 1998	1999 - 2003	2004 - 2008	2009 - 2014	1964 - 2014
Return	16.73%	15.42%	15.19%	0.77%	20.03%	15.21%
Standard Deviation	44.27%	46.52%	53.87%	44.03%	40.64%	41.31%
Skewness	0.27	0.32	0.35	0.27	0.18	0.38
Kurtosis	5.31	5.72	6.45	7.03	6.20	6.28
Max Drawdown	-28.14%	-28.84%	-34.83%	-32.99%	-27.08%	-29.78%
Serial Correlation	-9.03%	-9.42%	-6.10%	-5.65%	-5.48%	-2.10%
Sharpe Ratio	0.53	0.50	0.44	0.20	0.63	0.52
Certainty Equivalent	-22.1%	-26.7%	-42.6%	-39.4%	-12.1%	-18.3%
Daily abs(mean)/vol	5.55%	5.29%	5.18%	4.80%	5.43%	5.62%
Number of Stocks	3364	4587	3989	3626	2760	3184

For each statistic, we first take the mean over all stocks in each year, and then take the arithmetic mean over all years. The exceptions are skewness, kurtosis, and maximum drawdown, where we use the median in each year to avoid the effect of outliers. For the certainty equivalent, we assume a risk aversion coefficient of 3. *Number of stocks* is the number of stocks in the sample per year.

Note that taking the arithmetic mean across all years results in a large average annual return of 15.21%. Taking the geometric mean gives a more conservative number of 13.02%. However, since we use arithmetic averaging across years for the certainty equivalent, we also do it for returns to keep the comparison consistent.

1.8.7 Historical strategy performance and regime-switching

We want to link the historical performance of stop-loss strategies to the simulated results we obtained earlier for a MRS process. To do this, for every stock and every corresponding year of daily log returns, we estimate six MRS parameters:[8] the two means μ_1, μ_2, the two variances σ_1^2, σ_2^2, and the transition probabilities $P(1,1), P(2,2)$ of remaining in the same state. As before, we use the returns of the stop-loss strategy relative to the buy-and-hold strategy and control for time and size effects:

$$R_{i,t}^{S} - R_{i,t}^{BH} = \alpha + \beta_1 \mu_{1,i,t} + \beta_2 \mu_{2,i,t} + \beta_3 \sigma_{1,i,t}^2 + \beta_4 \sigma_{2,i,t}^2$$
$$+ \beta_5 P(1,1)_{i,t} + \beta_6 P(2,2)_{i,t} + \text{controls} + \epsilon.$$

We run the above regression for two strategies. Both are one-day strategies employing a 0% start level; however, one uses a stop level of −2%, while the other uses −5%. Note that since estimating MRS parameters on all stocks is computationally intensive, we randomly select 7% of all stocks in each year of study; this gives 11,341 stock-year observations across all years, which is still large. The results of the regression are listed in Table 1.7. The coefficients for the means in both regimes are negative and significant. This makes sense because a higher expected return leads to a worse performance on the strategy relative to the buy-and-hold strategy. We see that the coefficient for the mean in the bad regime is more negative and significant than in the good regime. This stems from the fact that the stop-loss strategy tends to earn its premium over the buy-and-hold strategy by staying out of the risky asset in bad times; therefore, its sensitivity to the expected return during those times is particularly large.

The relative return does not have strong dependence on the volatility in the good regime, probably because it is the bad regime that really matters for stop-loss strategies. On the other hand, the volatility in the bad regime has a large and negative coefficient. This is not consistent with our simulations results, where we predicted the stop-loss strategies to do better when this volatility is high since they would be able to switch to the bad regime "at the right time." On the other hand, we realize that very high volatility may lead the strategy to get back into the risky asset (and subsequently to keep trading in and out of it), so that transaction costs are high, while the benefits from staying in

[8] The estimation is performed in MATLAB using the MS_Regress package by Marcelo Perlin (2012).

Figure 1-8: Historical Stop-Loss Strategy Performance, 1964-2014; Risk-free Rate of Zero

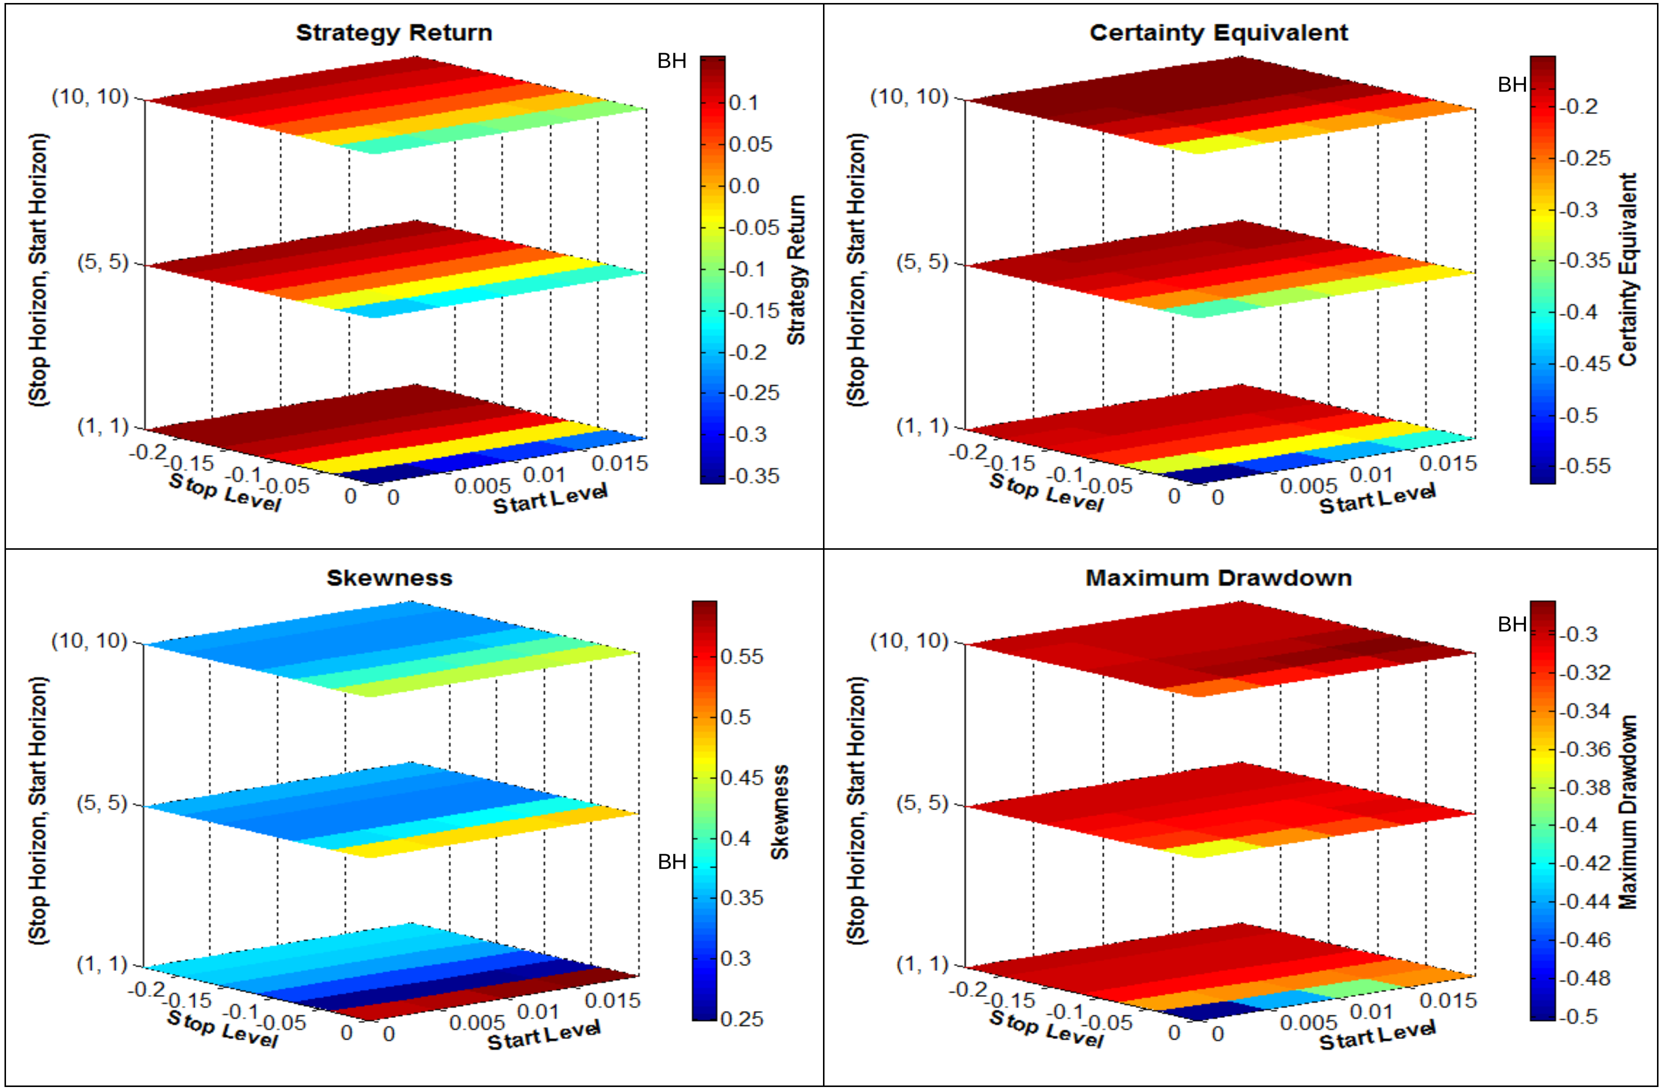

Historical performance of stop-loss strategies with varying levels of stop and start horizons and levels. The risk-free rate is set at 0%. All statistics are measured using annual returns and averaged out over all years in the sample. The performance statistic for the buy-and-hold is indicated on the colorbar with "BH."

the safe asset are limited.

Finally, we discuss the coefficients for the transition probability matrix. The coefficient for $P(1,1)$, the probability of staying in the good state, is negative because a higher value of this coefficient leads to a lower probability of switching to the bad regime, in which the stop-loss strategy does well relative to the buy-and-hold strategy. On the other hand, the coefficient for $P(2,2)$ is positive because a higher probability of staying in the bad state means that the bad regime is expected to persist for longer periods of time, leading to a greater potential for stop-loss outperformance. Both coefficients are significant at the 5% level only when the wider stop of -5% is used.

We conclude that the regression results are generally consistent with our simulations results. However, they are not as strong, and the coefficient for volatility is not consistent, possibly because the historical level of transaction costs has been high, while the MRS process does not always provide a good model for historical returns.

Table 1.7: Regressions of the Relative Return of Stop-loss Strategies on MRS Parameters of Stock Returns

Explanatory Variable	$\gamma = -2\%$	$\gamma = -5\%$	Average
Intercept	0.06%	0.06%	N/A
	(1.59)	(2.08)	
Mean 1	-0.004	-0.001	126.17%
	(-3.61)	(-1.51)	
Mean 2	-0.081	-0.039	-22.89%
	(-8.16)	(-5.65)	
Variance 1	0.009	-0.008	41.38%
	(0.74)	(-0.96)	
Variance 2	-0.452	-0.272	15.54%
	(-9.43)	(-8.21)	
P(1, 1)	-0.018	-0.035	75.48%
	(-0.86)	(-2.47)	
P(2, 2)	0.038	0.037	81.13%
	(1.75)	(2.52)	
Adj R-sq	23%	17%	N/A

In table 1.7 we regress the relative return of stop-loss strategies on MRS parameters fit to stock log returns. We control for time effect and firm size. In the first regression, we use the one-day strategy $\mathcal{S}(-2\%, 0\%, 1, 1)$ with stop level $\gamma = -2\%$. In the second regression, we use the strategy $\mathcal{S}(-5\%, 0\%, 1, 1)$ with stop level $\gamma = -5\%$.

1.8.8 Volatility-adjusting stop-loss strategies

So far we have considered stop-loss strategies where the stop level stays fixed throughout the whole investment horizon. With a fixed stop level, trading low volatility stocks results in the stop getting hit less frequently and, in turn, lower trading costs than for stocks with higher volatility. Therefore it may be worthwhile to consider a strategy with a stop-loss level that depends on the forecast volatility of the risky asset.

We use the exponentially weighted moving average (EWMA) model to produce volatility forecasts:

$$\hat{\sigma}_t^2 = \frac{1-\lambda}{1-\lambda^H} \sum_{i=1}^{H} \lambda^{i-1} r_{t-i}^2, \tag{41}$$

where λ is a parameter (the standard is 0.94, which is used in this paper), H is a time horizon (we use $H = 20$, corresponding to one month), and r_t are log returns.

Our volatility-adjusting strategy works as follows. At the end of each trading day we classify the stock as a high volatility or low volatility one. If the stock has high volatility, we use one pre-specified stop-loss level, and if the stock has low volatility we use another level. More specifically, we decide on a volatility threshold θ. At the end of day t, we calculate the forecast $\hat{\sigma}_t^2$ for stock volatility using equation (41). If the forecast is greater than the threshold θ, the stock is classified as a high volatility one and we set a stop level γ_2. Otherwise we set a stop level γ_1.

In Figure 1-9 we explore how changing the threshold θ and stop levels γ_1, γ_2 affects strategy performance. For convenience, we only consider two-week strategies (i.e. $I = J = 10$) with a start-gain level of 0%. The results are quite interesting. For returns it is best to use a wide stop both for high volatility and low volatility stocks. Among the two types of stocks, it appears that using a tighter stop hurts performance more for high volatility stocks than low volatility stocks. This is not surprising, because using a tight stop for a high volatility stock incurs very high transaction costs, causing poor returns. Overall, all strategies still produce a lower return (at most 14.3%) in comparison to the buy-and-hold strategy return of 15.2% per year.

When we look at the certainty equivalent (CE), the situation is different. Using a very wide stop is now not desirable, since often the volatility of the strategy remains high. This is particularly relevant for high volatility stocks – and we do indeed see that in Figure 1-9, where the highest CE levels are achieved when using a stop level of -10%

or -15% for such stocks. On the contrary, using a wide stop for low volatility stocks is still fine, because strategy volatility remains quite low, while performance is good. The best-performing strategies among the ones considered use a stop level of -15% to -25% for low volatility stocks, and either a stop level of -10% for high volatility stocks and a volatility threshold of 3% to 4%, or a stop level of -15% for high volatility stocks and a threshold of 2% to 4%. All these strategies have a CE between -14.4% and -14.3%, which is better than the buy-and-hold CE of -18.3%. It is also a small improvement over the CE of -14.9% for the best strategy among the "standard" stop-loss policies considered earlier (see Section 1.5.2).

The skewness of strategy returns is generally on par with the buy-and-hold strategy as long as we do not use tight stops for low volatility stocks. For the case when we do use tight stops, skewness improves quite significantly. However, this improvement is not due to a greater upside potential in the strategy, but rather by a large negative shift in its average return caused by high trading costs. We conclude that volatility-adjusting stops do not have a large effect on skewness. For maximum drawdown, we see a marginal improvement over the buy-and-hold strategy in almost all cases. When a stop of -5% is used for low volatility stocks, drawdown improves further, but still by a small amount.

We conclude that using stop-loss levels that depend on volatility forecasts can improve performance. For generating better returns, it is more important to use a wider stop for high volatility stocks than for low volatility ones. Overall, it is the stop-loss strategies with the widest stop levels that perform best, although still worse than the buy-and-hold strategy. If an investor cares about risk-adjusted returns, then a moderate stop level for high volatility stocks and a wide one for low volatility ones tends to generate the highest certainty equivalent, superior to the buy-and-hold strategy. Finally, stop-loss strategies tend to consistently improve maximum drawdown, although not by a lot.

Figure 1-9: Volatility-Adjusting Stop-Loss Strategy Performance, 1964-2014

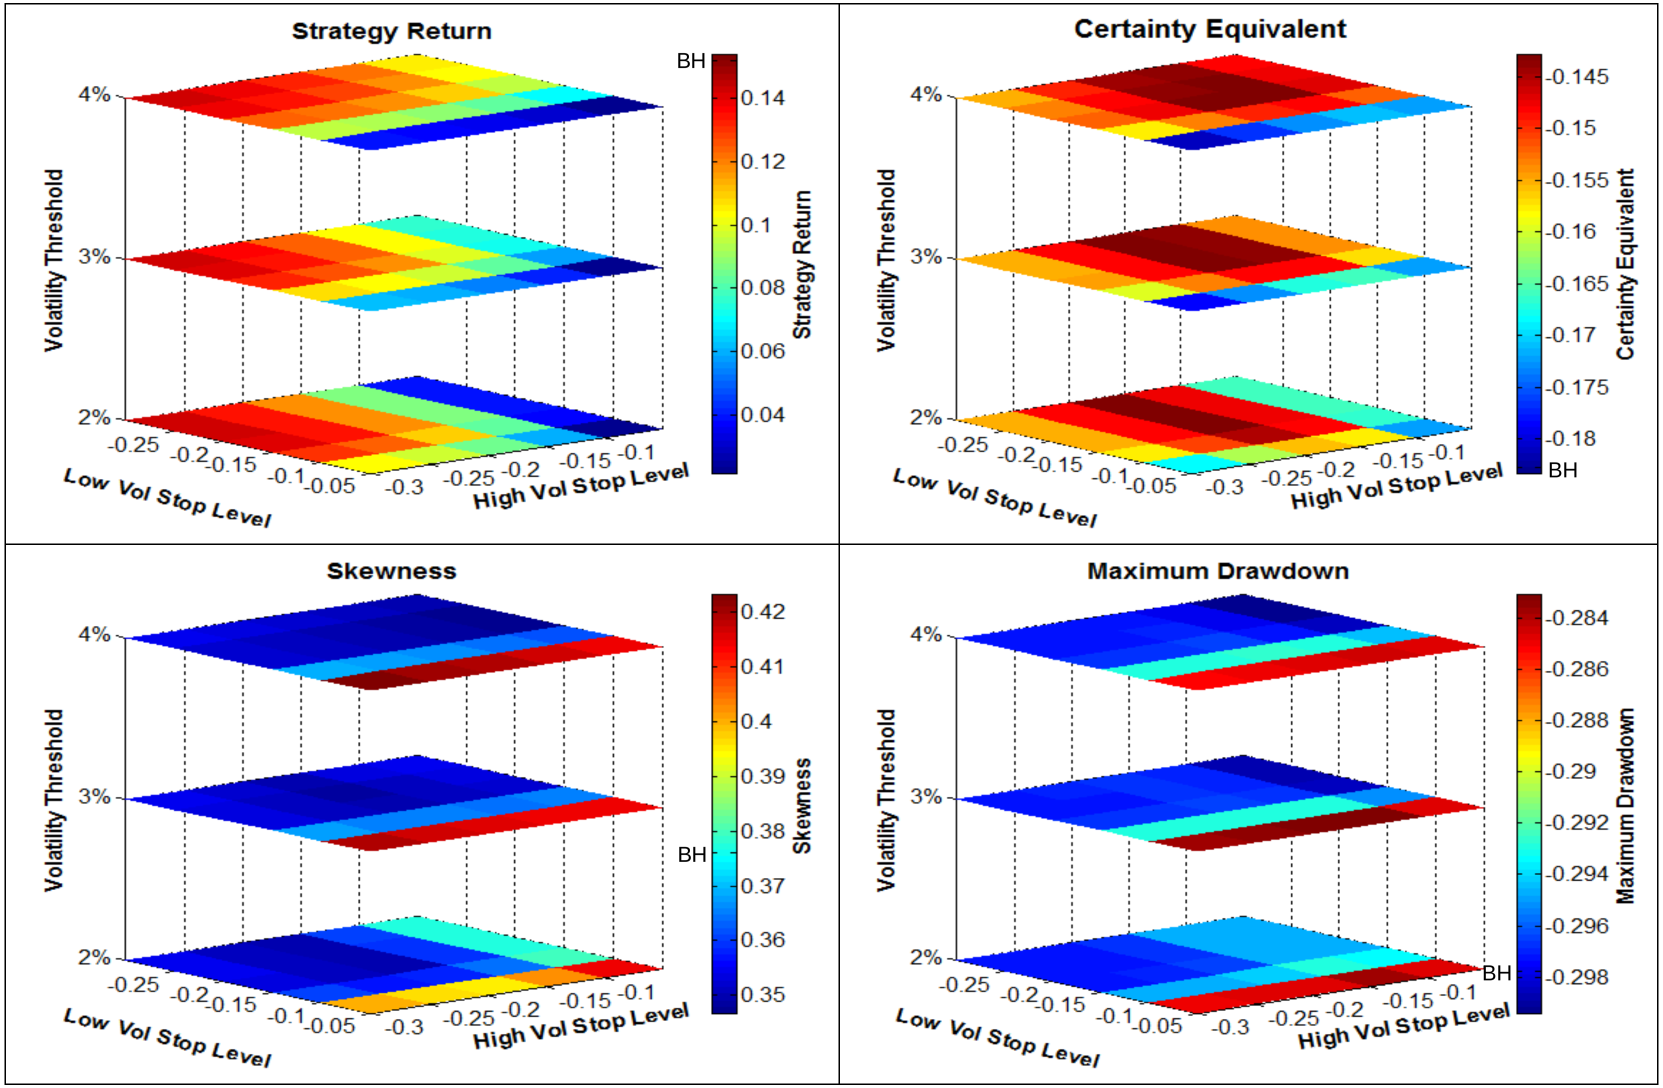

Historical performance of stop-loss strategies with volatility-adjusting stop levels. All strategies use a two-week period for past returns horizon (i.e., $I = J = 10$) and a start-gain level of 0%. We vary the volatility threshold θ, as well as the stop-loss levels γ_1, γ_2. The "safe" asset is the U.S. 30-day T-bill. All statistics are measured using annual returns and averaged out over all years in the sample. The performance statistic for the buy-and-hold is indicated on the colorbar with "BH."

Chapter 2

Measuring Risk Preferences and Asset-Allocation Decisions: A Global Survey Analysis

(joint work with Andrew W. Lo)

Abstract

We use a global survey of individual investors, financial advisors, and institutional investors to elicit their asset allocation behavior and risk preferences. We find drastically different behavior among these three groups of market participants. Most institutional investors exhibit highly contrarian reactions to past returns in their equity allocations. Financial advisors are also mostly contrarian; a few of them demonstrate passive behavior. Individual investors are, on average, extrapolative. To investigate further, we use a clustering algorithm to partition individuals into four distinct types: passive investors, risk avoiders, extrapolators, and everyone else. Across demographic categories, older investors tend to be more passive and more risk-averse, while more wealthy individuals are less risk-averse.

Keywords: Asset Allocation; Risk Aversion; Behavioral Finance; Retail Investors; Institutional Investors; Financial Advisors

JEL Classification: G02, G11, G23

2.1 Introduction

There are three major groups of participants in the investment management industry: individual investors, financial advisors, and institutions. Each of these groups has its own risk preferences and behavioral characteristics it uses in its investment decisions. We study the behavior of these groups by using the results of a comprehensive global survey covering 7,000 individuals, over 2,300 advisors, and over 600 institutional investors.

The breadth of our dataset sets it apart from earlier survey data in the literature. To the best of our knowledge, these surveys are the first to present the same set of questions to three distinct groups of market participants. This dataset covers 19 countries in the Individual Investor Survey and 16 countries in the Financial Advisor Survey. This global breadth provides us with insight into investment behavior by country, as well as allowing us to compare survey results across countries. Finally, all our survey subjects have a significant stake in the market: all the surveyed individual investors have a net worth above $200,000, while the financial advisors and the institutional investors are employed in the financial industry. As a result, their answers will generally be more realistic and have greater relevance for modeling investor behavior than the results of surveying students in a laboratory setting, as many other studies have done.

Our main goal is to understand how different market participants and different types of individuals compare along the dimensions of risk aversion and asset allocation. To this end, we poll members of these groups about their investment decisions under various historical and hypothetical scenarios. We obtain two sets of results. The first set of results shows that investors tend to be significantly more risk-averse and mostly extrapolative in their asset allocation, while institutions tend to be significantly less risk-averse and mostly contrarian in their investment decisions, with advisors falling in the middle of the risk aversion scale while also following a contrarian asset allocation strategy. The second set of results focuses on just individual investors—using a clustering algorithm applied to survey responses, we are able to identify four distinct types of investors: *passive investors*, *extrapolators*, *risk avoiders*, and *everyone else*. *Extrapolators* tend to decrease allocation in equities following bad market performance, and tend to increase allocation following good returns, extrapolating past trends. *Passive investors* leave their allocation unchanged in either scenario. *Risk avoiders* significantly cut their allocation to equities when they see large moves in the S&P 500 in either direction. The remaining investors is relegated to the last category, *everyone else*.

While the largest cluster of individuals in our dataset corresponds to passive investors, it also contains a significant number of risk avoiders and extrapolators. Evidence for each of these types is found in the literature, although most papers focus only on one type at a time. Agnew, Balduzzi, and Sunden (2003) and Dahlquist, Martinez, and Söderlind (2014) document that a large proportion of investors make no changes to allocations within their retirement portfolios over spans of several years. They note that this phenomenon may be linked to *inertia*, a widely recognized behavioral bias. Some major papers documenting extrapolation include De Bondt (1993), Greenwood and Shleifer (2014) using survey data, as well as Benartzi (2001) and Choi et al. (2009) which look at historical 401(k) account allocations. Finally, Ben-David and Hirshleifer (2012) investigate individual trading records and derive a V-shaped probability distribution of selling a stock as a function of profit. This in part may be driven by the cluster of risk avoiders identified from our survey.

In comparison, we find that most financial advisors and institutional investors are contrarian in allocation strategy—that is, they would change equity allocation in the direction opposite to recent returns on the S&P 500. This contrasts with the overall behavior of individual investors, who on average are extrapolators.[1] The differences in the reactions across these three groups of market participants are significant and very large. We note that a few earlier studies have viewed individuals as momentum traders, and institutions as contrarians.[2] However, these studies consider shorter-term horizons than ours, and focus on trading behavior. Our survey asks about asset allocation, a strategic and longer-term investment decision, rather than short-term trading, which potentially could be affected by excessive speculation on the part of individual traders, or by liquidity considerations of institutions. A recent paper by Haan and Kakes (2011) does focus on asset allocation of Dutch institutional investors, and concludes that they tend to be contrarians.

Our results have another important implication, one that arises from the differences in response between financial advisors and individual investors. We find that advisors

[1]While a large number of individual investors are passive or risk avoiders, both of these groups have symmetric reactions to large moves in the S&P 500: they either do nothing, or they significantly decrease equity allocation.

[2]See, for example, Griffin, Harris, and Topaloglu (2003), Jackson (2003), and Kaniel, Saar, and Titman (2008).

generally advise their clients to change their allocation in the opposite direction of the typical preference of the individual investor. It may be that advisors recognize the excessive tendency of investors toward extrapolation, and try to mitigate this effect by giving "contrarian advice". Also, the proportion of advisors who suggest a significant decrease in equity allocation when seeing large S&P 500 moves is much smaller than the proportion of individual investors who would implement such a change. As a result, advisors may also provide the significant benefit of ensuring their investors stay invested in the markets despite periods of high volatility, and hence earn higher returns in the long run.[3] Overall, our findings suggest that financial advisors are of direct benefit to most individual investors.

We also investigate how individual investor demographics can be used to predict their anticipated investment decisions, particularly in response to large shocks. While very limited data on demographics and past trading behavior is available from the survey, we are able to obtain a significant improvement in prediction accuracy over the baseline by using standard analytics approaches, particularly the random forest model.

Finally, we compare risk aversion across the three groups, as well as within investor demographic categories. Individual investors are significantly more risk-averse than financial advisors, who are in turn more risk-averse than institutional investors. Individual risk aversion increases with age, and decreases with wealth. This is consistent with previous literature linking risk aversion to age, wealth, and education; see Masters (1989), Pålsson (1996), and Hartog, Ferrer-i-Carbonell, and Junker (2002).

In Section 2.2 we outline the survey methodology and the estimation of risk aversion. We compare survey responses across individuals, advisors, and institutions in Section 2.3. We focus on individual investors in more detail in Section 2.4 and we conclude in Section 2.6.

2.2 Methodology

We use data from three separate but closely related surveys: the Natixis Global Survey of Individual Investors, the Natixis Global Survey of Financial Advisors, and the Natixis Global Survey of Institutional Investors. Each survey involved two sets of questions. The

[3]Winchester, Huston, and Finke (2011) find that investors using a financial advisor are 1.5 times more likely to stick to long-term investment decisions.

first set originated with Natixis Asset Management for their own research purposes. The second set was created by us, in coordination with Natixis, for studying the behavioral aspects of investor decision-making.

Survey Questions

We asked four questions in each survey. The first involved preferences among potential gambles, and was used to elicit the risk aversion coefficient among respondents. The second and third asked how respondents would change their investment allocation as a result of large negative or positive moves in the S&P 500. Possible responses were "significantly decrease equity allocation", "slightly decrease equity allocation", "do nothing", "slightly increase equity allocation", and "significantly increase equity allocation". In the case of financial advisors, we asked them how they would advise their clients to act in such situations. The final question asked if and when investors decreased their allocation during the 2007–2009 Financial Crisis. The exact formulation of the questions is included in the Appendix.

Natixis Asset Management commissioned CoreData Research to conduct each survey via an online questionnaire. The Individual Investor Survey was carried out in March 2015 and involved 7,000 individuals from 17 countries. Each investor needed to have a minimum net worth of $200,000 (or Purchasing Power Parity equivalent) to participate. The Financial Advisor Survey was conducted in June–July 2015 and involved 2,400 advisors from 16 countries. Since some advisors opted to not complete the behavioral part of the survey, we had a total of 2,342 advisor observations. The Institutional Investor Survey was carried out in October 2015 and involved 660 respondents from 29 countries. All respondents had to be decision makers working in the institutional investment industry, such as Chief Investment Officers, pension fund managers, and investment portfolio managers.

Table 2.1 provides summary statistics across the three groups on age, gender, and net worth/assets. In the Appendix we also include a breakdown of respondents by country in the individual and advisor surveys, as well as the list of institution types covered in the institutional survey.

It is important to keep in mind that the surveys give us information about what investors *think they would do* under various market scenarios. Thus the results pertain to

<table>
<tr><th>Survey Group</th><th># Subjects</th><th colspan="3">Summary</th></tr>
<tr><td rowspan="2">Individual Investors</td><td rowspan="2">7000</td><th>Age</th><th>Gender</th><th>Net Worth</th></tr>
<tr><td>Generation X: 32%
Generation Y: 37%
Baby Boomers: 28%
Pre-Baby Boomers: 3%</td><td>Male: 59%
Female: 41%</td><td>Mass Market: 21%
Mass Affluent: 24%
Emerging HNW: 25%
High Net Worth: 30%</td></tr>
<tr><td rowspan="2">Financial Advisors</td><td rowspan="2">2342</td><th colspan="3">Advisor Characteristics</th></tr>
<tr><td colspan="3">Average age: 46 years
Male: 79%; Female: 21%
Average personal book of business: $28.3 million</td></tr>
<tr><td rowspan="2">Institutional Investors</td><td rowspan="2">660</td><th colspan="3">Organization Assets</th></tr>
<tr><td colspan="3">Less than $2 billion: 18%
$2 billion – $5 billion: 25%
More than $5 billion: 57%</td></tr>
</table>

Table 2.1: Summary statistics for the three survey groups on age, gender, and net worth/assets. Different groups were presented with different demographic questions, and the summary statistics are not directly comparable across groups. Gender and age were not asked for institutional investors.
The definitions of generations of investors and their net worth classifications are in the Appendix. *HNW* stands for High Net Worth. *Organization Assets* means the size of assets for which the respondent's organization is responsible.

investors' conditional expected changes to their portfolios; to obtain actual changes, we would need to use historical data on portfolio allocations or on investor trades. However, we obtain highly significant differences across various investor groups, so that even in the presence of potential noise in self-reporting how investors anticipate their future decisions, we may still draw reliable conclusions from these comparisons.

Estimating Risk Aversion

We estimate the risk aversion of survey respondents using the popular technique introduced by Holt and Laury (2002), and later modified by Dave et. al (2010). The idea by Holt and Laury is to present subjects with pairs of gambles and ask which gamble they would choose within each pair. Dave et al. extend this idea by presenting a list of six different gambles and asking subjects to pick one. Table 2.2 lists the six gambles used in the Natixis surveys, with gamble 1 being the safest and gamble 6 being the riskiest. Note that these are the same gambles as in the Dave et al. study, except they are of greater magnitude, given that our subjects are professional investors or industry professionals, who tend to care about relatively large bets when investing.

The choice of one gamble among six gambles is converted to a choice of one gamble in each of five different pairs of gambles: gamble 1 vs. gamble 2, gamble 2 vs. gamble 3, ..., gamble 5 vs. gamble 6. The conversion is done as follows. If a person chose a particular gamble k among the six gambles, then he would choose gamble $i+1$ over gamble i as long as $i+1 \leq k$, and would choose gamble i over gamble $i+1$ as long as $i \geq k$. This is valid because gambles $1, 2, \ldots, 6$ are increasing in risk, in that order.

Gamble	**Outcome 1**		**Outcome 2**	
	Probability	Payoff	Probability	Payoff
Gamble 1	50%	$28,000	50%	$28,000
Gamble 2	50%	$36,000	50%	$24,000
Gamble 3	50%	$44,000	50%	$20,000
Gamble 4	50%	$52,000	50%	$16,000
Gamble 5	50%	$60,000	50%	$12,000
Gamble 6	50%	$70,000	50%	$2,000

Table 2.2: List of six gambles presented to survey participants. The subjects were asked to choose which one of the gambles they would prefer. Each gamble involves two outcomes, each of which has a 50% probability of occurring. The first gamble can be viewed as a “sure" outcome.

We assume that respondents have constant relative risk aversion (CRRA) utility. The risk aversion coefficient r will be used as a proxy for risk aversion among different groups throughout the paper. When considering a particular gamble i, a subject first evaluates their expected utility:

$$E(U_i) = p_{i,1}U(x_{i,1}) + p_{i,2}U(x_{i,2}); \; U(x) = \frac{x^{1-r} - 1}{1 - r} \tag{1}$$

where $x_{i,1}, x_{i,2}$ are the two possible payoffs for the gamble, and $p_{i,1}, p_{i,2}$ are their associated probabilities of occurring. Then, for each pair of gambles i, j, after calculating expected utility, the subject picks gamble i with probability:

$$\mathbb{P}(\text{ choose gamble } i) = \frac{[E(U_i)]^{1/\mu}}{[E(U_i)]^{1/\mu} + [E(U_j)]^{1/\mu}} \tag{2}$$

where μ is a noise parameter, since a subject may actually pick gamble j even if gamble i has higher expected utility. Values close to 0 signify little deviation from expected utility theory.

Using equations (1), (2), and survey responses, we calculate the likelihood function for parameters μ, r. We then estimate parameters using the maximum likelihood method, and use the Hessian to estimate standard errors. The risk aversion estimates are compared across the three groups of subjects as well as across individual investor demographic categories.

2.3 Investors, Advisors, and Institutions

In this section, we present our comparison of allocation decisions and risk aversion across investors, advisors, and institutions. The results are striking. Investors appear to be mostly extrapolative in their changes in equity allocation, while advisors and institutions are predominantly contrarian. The proportion of "passive" respondents is much higher for individuals and financial advisors than for institutional investors. Finally, investors are more risk-averse than advisors, who are in turn more risk-averse than institutions.

Asset Allocation Decisions

All individual and institutional investors were asked two questions pertaining to investment decisions. The first asked how they would change their allocation to equities if the S&P 500 declined by 10–20% during the next six months, and other assets performed as expected. The second asked the same question in a scenario where the S&P increased by 10–20%. Financial advisors were asked two similar questions about how they would *advise their clients* to change equity allocation. There were five possible responses: "large decrease", "slight decrease", "do nothing", "slight increase", and "large increase".

Figures 2-1 and 2-2 plot the distributions of responses for the two questions. In the scenario of an S&P 500 fall, 48% of individuals would decrease equity allocation,[4] in comparison to just 17% for advisors and 20% for institutions. At the same time, 71% of institutions and 53% of advisors would increase equity allocation, a much higher proportion than the corresponding 18% for individuals. For a rise in the S&P 500, the results are basically reversed. For individuals, 32% would decrease allocation and 36% would increase it, whereas the corresponding numbers are 49% and 13% for advisors, and 66% and 13% for institutions. Thus, on average, individual investors would change their allocation in the same direction as a recent S&P 500 move, while financial advisors and

[4] The proportion of respondents decreasing equity allocation is calculated by counting all respondents who answered with "large decrease" or "slight decrease" to the survey question.

institutional investors would change allocation in the opposite direction.

We obtain more insight into these differences by looking at the response distributions in more detail. One aspect that stands out is the "extreme contrarian" responses of institutional investors. 36% would significantly increase allocation following a fall in the S&P 500, whereas 27% would significantly decrease allocation following a rise in the S&P 500. This is much higher than the corresponding 4% and 17% for individuals, or the 13% and 10% for advisors. While advisors also give contrarian responses, their predominant response it to change allocation only *slightly*: 40% of advisors would favor a slight increase after a fall in the S&P 500, and 39% of advisors would favor a slight decrease after a rise in the S&P 500—much higher than the comparable rates among individual investors. Finally, we notice that while investors on average extrapolate, there is a significant fraction in both scenarios that prefers to significantly decrease equity allocation: 21% for an S&P 500 fall and 17% for an S&P 500 rise. We will later show, using clustering techniques on the individual investor dataset, that these numbers are mainly driven by the "risk avoider" class of investors.

As a robustness check, we also compare the responses across individuals and advisors for each country included in both the Individual Investor and the Financial Advisor Surveys.[5] The results are shown in Figures 2-3 and 2-4. We again see that there is a much larger proportion of individual investors who would decrease equity allocation after seeing an S&P 500 fall than there is of advisors; at the same time, many more advisors would increase allocation in this scenario compared to individuals. The differences are significant for each country at the 1% level. In the scenario of an S&P 500 rise, the results are reversed; this time, advisors are the ones who are more likely to decrease equity allocation. It is important to note that the differences between advisors and individual investors who decrease equity allocation are now smaller (although still significant at the 1% level for all but three countries). This again is caused by "risk avoider" investors who significantly decrease equity allocation following a rise in the S&P 500.

The fifth category of response in these two scenarios is the "do nothing" response. Its levels for individual investors and financial advisors are quite similar in both scenarios, and relatively large: 34% and 30% in case of an S&P 500 fall, and 32% and 38% in case of an S&P 500 rise. In contrast, only 9% of institutional investors would do nothing

[5] We do not include institutional investors in the comparison because the sample sizes at the country level are small in the Institutional Investor Survey.

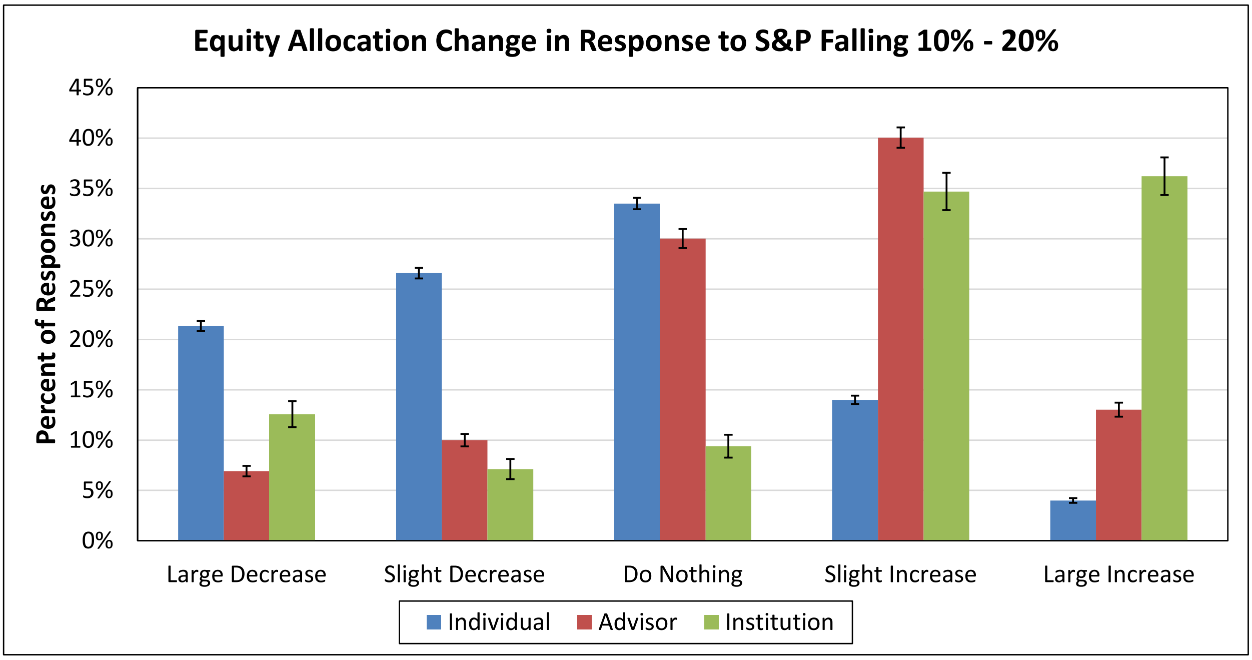

Figure 2-1: Reactions to a decrease in the S&P 500 across three groups. For each group and each possible answer, we show error bars corresponding to one standard error calculated assuming each respondent chooses either that particular answer, or any other answer.

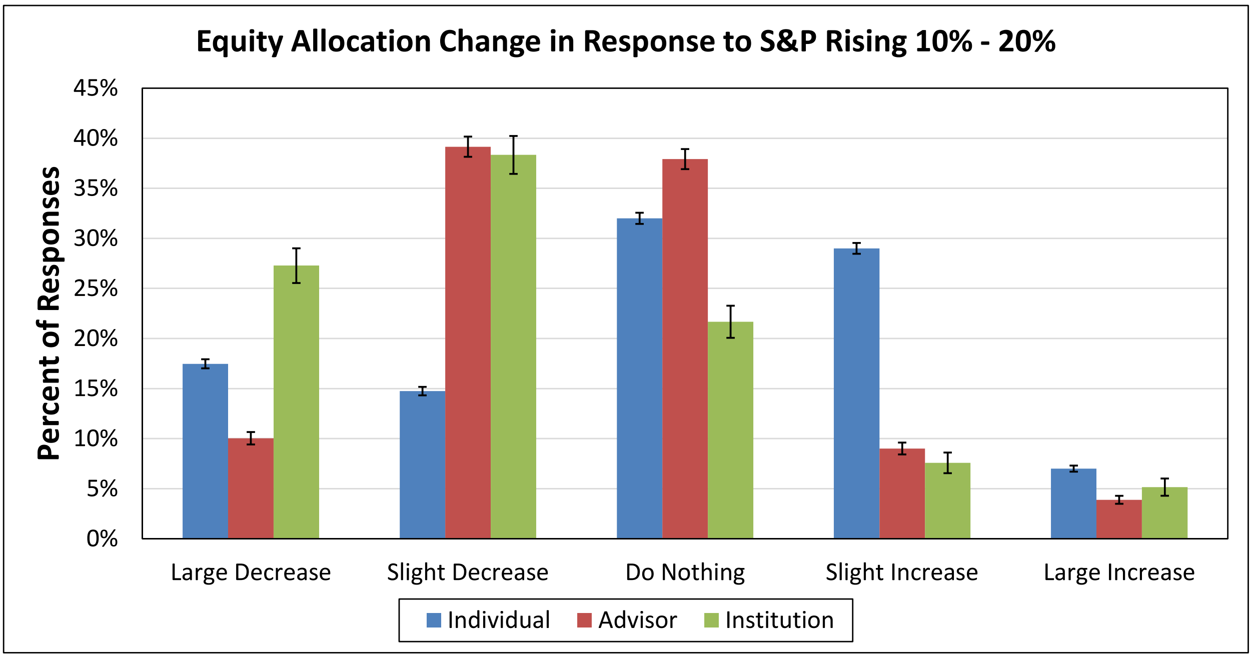

Figure 2-2: Reactions to an increase in the S&P 500 across three groups.

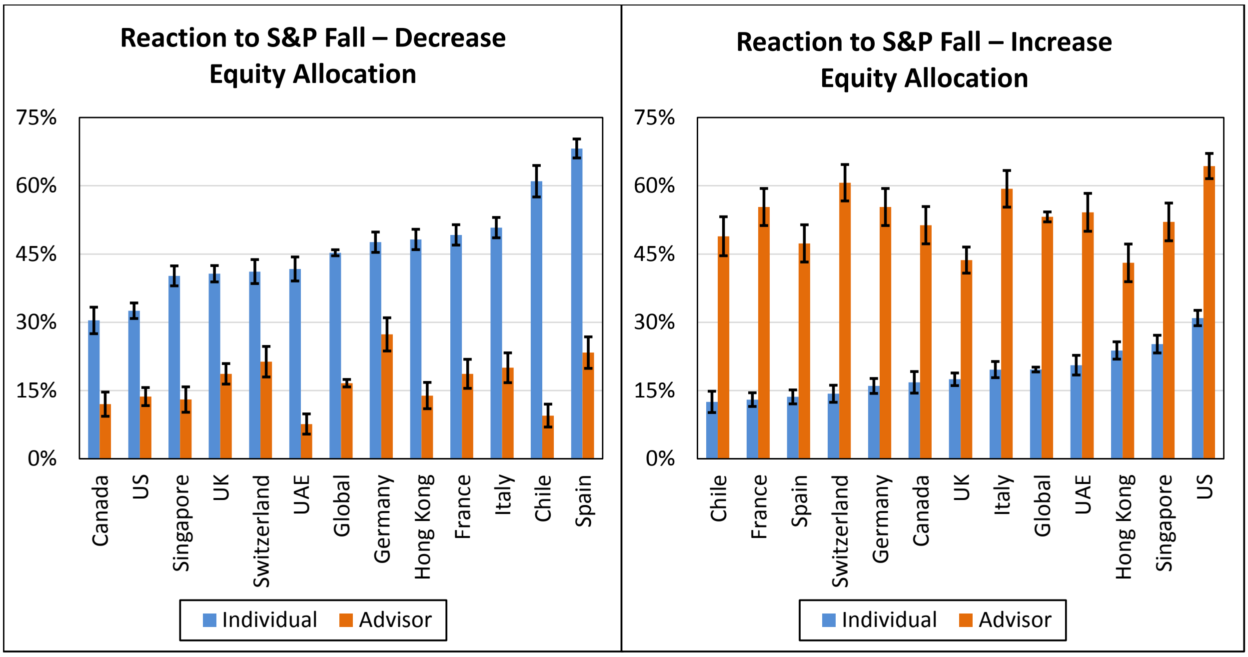

Figure 2-3: Reactions to a decrease in the S&P 500 across individual investors and financial advisors, split up by country. Error bars correspond to one standard error.

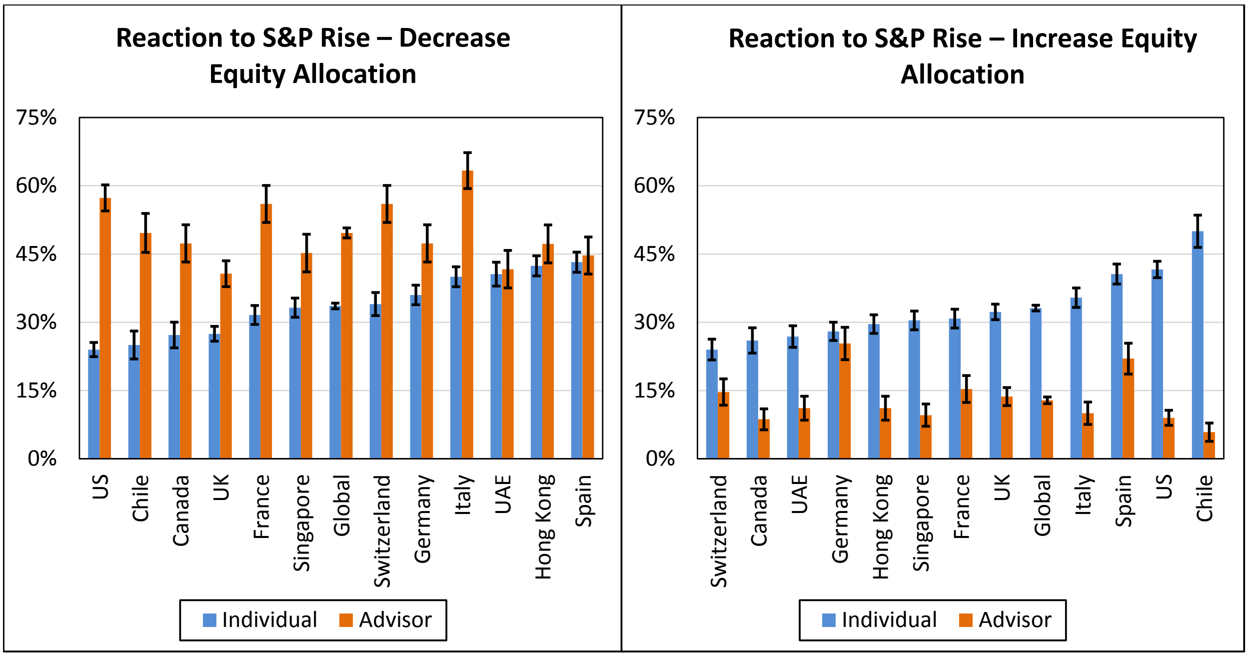

Figure 2-4: Reactions to an increase in the S&P 500 across individual investors and financial advisors, split up by country.

following a decrease in the S&P 500; this number is 22% following an increase in the S&P. The more passive responses for individuals and advisors may be explained by the fact that many have long-term investment objectives that do not require significant changes to their asset mix. As a result, they would generally not be affected by large fluctuations in asset prices. In contrast, institutional investors often have shorter investment horizons, and their performance may be evaluated at shorter frequencies. Because of this, they may need to react to changes in the relative prices of different asset classes more often.

The observed differences in active response between groups have several potential explanations. The apparent tendency of a large number of individual investors toward extrapolative allocation may be an inherent aspect of their behavior. Our survey is insufficient to understand why this behavior comes about. However, we note that pervasive evidence for similar behavior has been documented in other studies; see Greenwood and Shleifer (2014) for a comprehensive survey. Financial advisors may recognize that some of their individual investor clients excessively extrapolate, and instead advise them to apply a contrarian allocation strategy, as observed in the survey. Also, should the S&P 500 decline significantly, an advisor may view that circumstance as a good entry point for a client with a long investment horizon.

The contrarian behavior of financial advisors may also be explained by long-term investment objectives, which are typically planned to maintain a target asset mix over several years. For example, if equities move significantly relative to bonds over the short term, then client allocation will experience a large deviation from its target mix, and advisors may propose a contrarian reallocation to return it to the target.

The target-mix story may also explain the extreme contrarian response of institutions, especially if their performance is evaluated relative to a benchmark. However, because a large proportion of institutional respondents would *significantly* decrease or *significantly* increase allocation, there are probably more factors at play. It is possible that some institutional investors are employing value strategies or are engaged in distressed investing, which results in contrarian trading when asset class prices deviate from their earlier relationship. Another possibility may be that some investors (e.g. pension plans) have a target return in mind. If recent performance is very good, they may cut portfolio risk by moving out of equities in order to have a safer portfolio for the rest of the year, while likely still hitting their target.

Overall, because individual investors tend to have extrapolative reactions, while institutional investors usually have contrarian ones, we conclude that institutions generally take the other side of individual investor trades in broad asset allocations. Note that Greenwood and Shleifer (2014) propose firms may be also involved in accommodating individual investor demand through equity issuance.

Risk Aversion

We next compare risk aversion across the three survey groups. Figure 2-5 shows the distribution of preferences for the six gambles across respondents. Recall that gamble 1 is safest, and gamble 6 is riskiest. We see that 40% of individuals choose gamble 1, much higher than 27% for advisors, which in turn is higher than 17% for institutions; the differences are significant at the 1% level. At the other end of the spectrum, a significantly higher proportion of institutions choose gambles 5 and 6 in comparison to individuals and advisors. These observations strongly suggest that individual investors are the most risk-averse of the three groups, while institutional investors are the least risk-averse. Note that we are able to make this conclusion from the distribution of responses alone, without making assumptions about the utility functions of respondents.

To investigate further, we add the assumption that all subjects have CRRA preferences, and carry out the estimation procedure for the risk aversion coefficient as discussed in the methodology. The results are shown in Figure 2-6. Individual investors have the highest risk aversion coefficient at 1.14, followed by financial advisors at 0.85 and institutional investors at 0.39. The pairwise differences in coefficient estimates are very large, and are again significant at the 1% level. Our results are consistent with the general intuition that individuals are generally the most risk-averse group of market participants, while institutional investors are the least risk-averse one. Likewise, it is plausible for financial advisors to fall in the middle of the spectrum, given the fact that, while they do work in the investment management industry, they generally do not directly manage money for their clients and so do not take outright bets in the markets, whereas institutions do.

We have presented evidence that individual investors are, on average, extrapolators, while financial advisors and institutional investors are contrarians. Contrarian behavior is particularly strong for institutions. Furthermore, institutions would usually reallocate portfolios more actively in response to large S&P 500 moves, while a significant proportion of individuals and advisors would do nothing in those situations. Individuals are by far more risk-averse than financial advisors, and advisors are much more risk-averse than institutional investors.

2.4 Individual Investor Decisions

We have seen that a large proportion of individual investors tend to extrapolate market performance when making their equity allocation decisions. However, we have also observed that a few of them tend to consistently decrease equity allocation. This suggests that we cannot label all investors as extrapolators, and it makes sense to study the individual dataset in more detail. We start by looking at the dependence of investor risk aversion and preference for active investing on the demographic factors of age, gender, and net worth. The strongest results are observed for age, older investors tending to be more risk-averse and also more passive. We then run a clustering algorithm to partition investors into four groups: passive investors, risk avoiders, extrapolators, and everyone else. We also compare the demographic breakdowns of the different investor types.

Asset Allocation

We used three questions from the survey to elicit the degree of investor passivity. The first two questions asked about their allocation under different S&P 500 moves, as discussed in the last section. The third question asked when during the Financial Crisis of 2007-2009 investors decreased their equity allocation. Possible answers ranged from “the second half of 2007" to “the second half of 2009". We also included a “do nothing" response for investors who did not significantly decrease allocation. Figure 2-7 compares the proportions of investors choosing the passive response for each of the three questions across the different demographic categories. Note that we do not have exact numbers for age or net worth, but rather a classification into one of four possible groups, listed in Table A.5 in the Appendix.

Within each demographic category, the percent of respondents who choose to do nothing is similar in all three scenarios, the percentage responding passively to the Financial Crisis only slightly higher than the same response to the hypothetical moves in the S&P

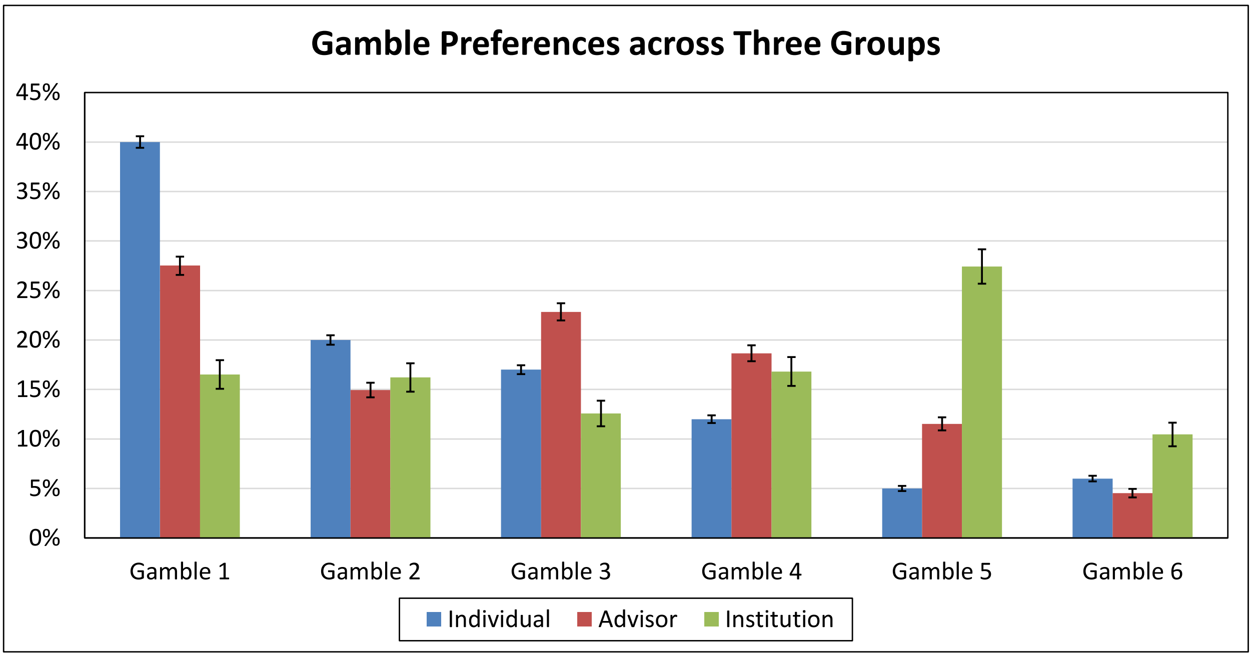

Figure 2-5: Distributions of gamble preferences across three groups. For each group and each possible answer, we show error bars corresponding to one standard error calculated assuming each respondent chooses either that particular gamble, or any other gamble.

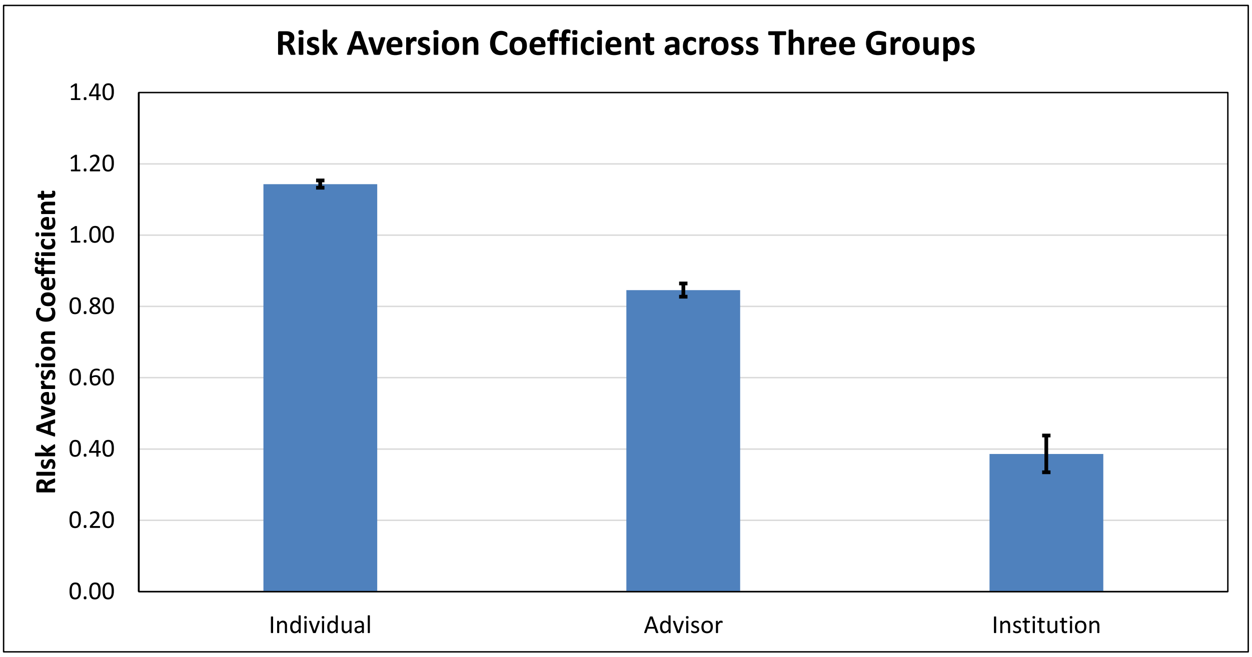

Figure 2-6: Estimated Risk Aversion coefficients across the three groups. Error bars correspond to one standard error.

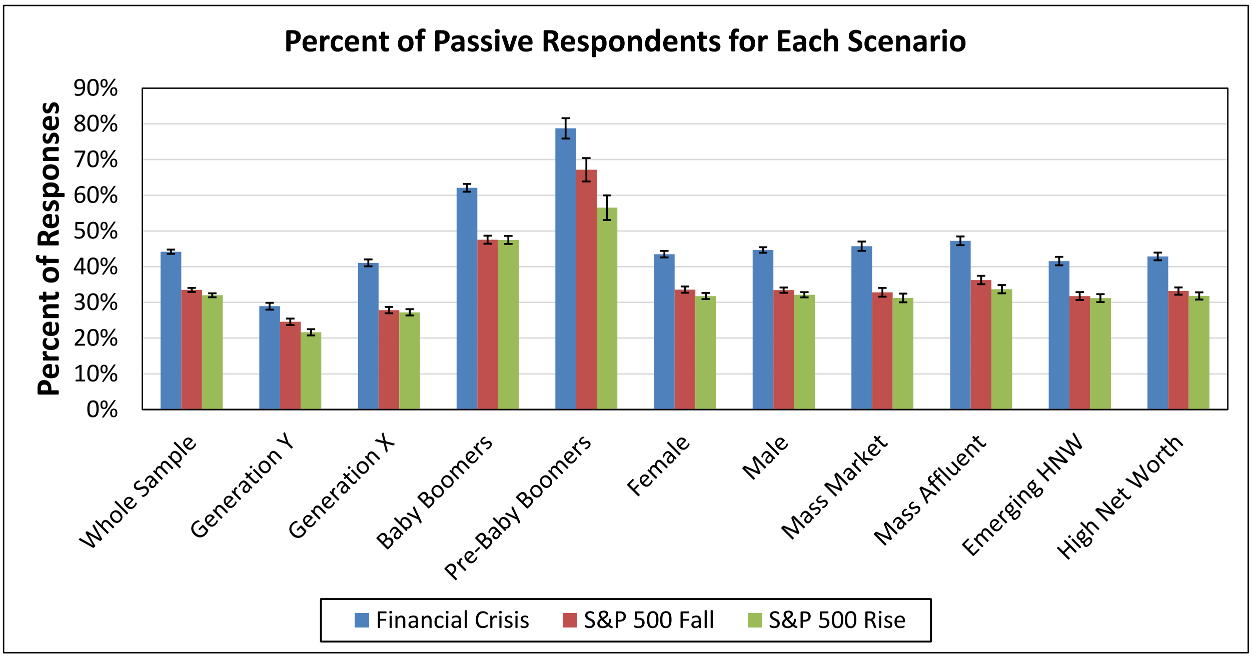

Figure 2-7: Percent of passive respondents for each scenario across different individual investor demographic categories. The definitions of investor demographic categories are in Table A.5 in the Appendix. *HNW* stands for High Net Worth. Error bars correspond to one standard error, assuming the respondent chose either the "passive" response or any other response.

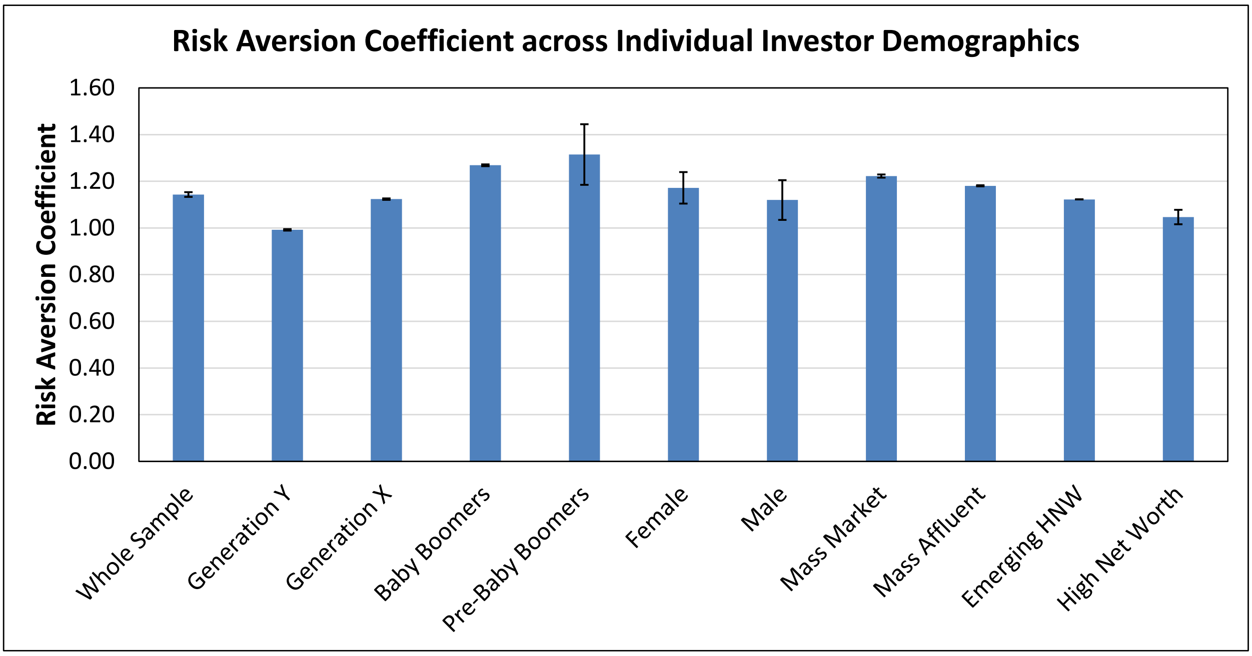

Figure 2-8: Estimated Risk Aversion coefficients across individual investor demographic categories. Error bars correspond to one standard error.

500. However, there are large and statistically significant (at 1%) pairwise differences across age categories, where older investors are much more likely to be passive. Comparing the different ends of the spectrum, 79% of pre-Baby Boomers did not significantly decrease equity allocation around the time of the crisis, while this number is just 29% for Generation Y investors. The greater observed percentage of older investors being more passive may be a result of their inherent behavior; this is consistent with other empirical evidence, for example, by Dorn and Huberman (2005) and Grinblatt and Keloharju (2009). Another potential explanation is that younger investors generally tend to have a higher allocation of equities in their portfolios, and therefore would react more to a large change in equity prices. The differences across gender and net worth categories are not large, and almost in all cases not significant at the 5% level.

Risk Aversion

Figure 2-8 compares the estimated risk aversion coefficient across demographic categories. Risk aversion increases with age and decreases with net worth; the pairwise differences across generations and across net worth categories are significant at the 1% level (except for pre-Baby Boomers; standard errors for this cohort are large, in part, due to the small sample of 207 respondents). These results have been documented earlier in multiple studies, including Pålsson (1996) for age and Cohn et. al (1975) for net worth. Women appear to be slightly more risk-averse than men, although the difference is not significant at the 5% level. We also include a comparison of individual risk aversion across countries in Figure 2.7.1 in the Appendix. Standard errors are large for a few countries due to the small sample size, and large differences between countries are generally not seen. The only outlier is Hong Kong, where investors are significantly less risk-averse than in almost every other country considered (the differences are significant at the 5% level).

Cluster Analysis

We now perform clustering on the individual dataset of 7,000 investors. We use the responses to the three asset allocation questions described earlier. The verbal responses are transformed into numbers as shown in table 2.3; the numbers range from 1 to 6 for the reaction to the Financial Crisis, and from -2 to 2 for the two questions on reactions to S&P 500 moves. We then run a k-means clustering algorithm assuming four clusters. We choose four clusters since this partition seems to give the best results (i.e., the most distinct and "cleanest" clusters) compared to our trials using 2, 3, 5, and 6 clusters. For robustness we ran the algorithm for different values of the random seed, and performed hierarchical clustering on the dataset; the results were generally the same in all cases.

Question	**Response Coding**
Reaction to Financial Crisis	Timing of Decrease in Equity Allocation: 1 - Second Half of 2007 2 - First Half of 2008 3 - Second Half of 2008 4 - First Half of 2009 5 - Second Half of 2009 6 - No Significant Decrease
Reaction to S&P 500 Fall	Change in Equity Allocation: -2 - Significant Decrease -1 - Slight Decrease 0 - Do Nothing 1 - Slight Increase 2 - Significant Increase
Reaction to S&P 500 Rise	Change in Equity Allocation: Same as for Reaction to S&P 500 Fall

Table 2.3: Coding of investor responses for the clustering algorithm. The exact formulation of the questions and possible responses is in the Appendix.

Table 2.4 shows the results of the clustering procedure. Cluster 1 consists predominantly of *passive investors*, since the average responses are "close" to the response corresponding to not changing equity allocation in each of the three scenarios: 5.9 for allocation around the crisis, -0.3 for an S&P 500 fall and 0.0 for an S&P 500 rise. Cluster 2 contains mostly *risk avoiders* who significantly cut allocation following large moves in the S&P 500: average responses are -1.4 for an S&P 500 fall and -1.6 for an S&P 500 rise. Cluster 3 is composed of *extrapolators* who tend to change equity allocation in the same direction as a recent S&P 500 move: average responses are -1.5 for an S&P 500 fall and 1.0 for an S&P 500 rise. Finally, cluster 4 contains *everyone else.*

To further validate our clustering approach, we look at the distributions of responses within each cluster, shown in Figure 2.7.2 in the Appendix. It is evident that the distributions for both the risk avoider and the extrapolator clusters are tightly clustered around the corresponding means for the reactions to S&P 500 moves. The same is true for the passive cluster for all three questions (and especially, for the allocation around the Financial Crisis). The response distribution of the "everyone else" cluster is quite spread out for all three responses, except for the response to an S&P 500 fall, explaining why we do not label the investors in this cluster as following any well-defined type of behavior. We conclude that the clustering into four groups is indeed successful, and the descriptions of the clusters are appropriate.

Clustering Analysis of Individual Investors				
	Cluster 1	**Cluster 2**	**Cluster 3**	**Cluster 4**
% Respondents	51%	16%	14%	19%
	Allocation Decisions			
Reaction to Crisis	5.9	2.7	2.5	2.6
Reaction to S&P 500 Fall	−0.3	−1.4	−1.5	0.7
Reaction to S&P 500 Rise	0	−1.6	1	0.2
	Demographics			
Gender	0.41	0.42	0.41	0.41
Age	1.13***	0.69***	0.86	0.8
Net Worth	1.64	1.62	1.63	1.72*
Advised	0.56	0.57	0.67*	0.63
Satisfied with 2014 Ret.	0.59	0.44***	0.55	0.61
Retired	0.09***	0.02	0.02	0.03
Gamble Preference	1.19***	1.64	1.47	1.62
Risk Aversion Coefficient	1.24***	0.98	1.08	1.01

Table 2.4: Clustering of allocation decision responses from the Individual Investor Survey. For each cluster, we present the percent of respondents and the mean response based on the response coding in Table 2.3.
We also list the mean values of demographic categories across clusters. For *Gender*, *Male* = 0, *Female* = 1. For *Age*, *Generation Y* = 0, *Generation X* = 1, *Baby Boomers* = 2, *Pre-Baby Boomers* = 3. For *Net Worth*, *Mass Market* = 0, *Mass Affluent* = 1, *Emerging HNW* = 2, *High Net Worth* = 3. The definitions of demographic categories are in Table A.5 in the Appendix. *Advised* is an indicator for if an investor uses a financial advisor. *Satisfied with 2014 Ret.* is an indicator for if an investor was satisfied with their 2014 investment returns. *Retired* is an indicator for if an investor is retired. *Gamble Preference* corresponds to the one of six gambles from Table 2.2 chosen by the investor; if an investor chooses gamble i, his response is recorded as $i-1$ for calculating the mean, so that responses range from 0 to 5. *Risk Aversion Coefficient* is the estimated risk aversion coefficient based on the responses in each cluster.
For each category, we color green the cell corresponding to the cluster with the highest mean value. We test for how significant the difference is between the highest mean and second-highest mean across the clusters; the result of the test is reported in terms of number of stars in the cell. * means significance at the 5% level, *** means significance at the 0.1% level; no stars means no significance at the 5% level. We color red the cell corresponding to the cluster with the lowest mean value, and perform the same test comparing the lowest mean and second-lowest mean across the clusters.

It is important to note the relative sizes of the clusters from Table 2.4. Passive investors are the largest cluster and make up 51% of the whole sample. Risk avoiders consist of 16% of all investors, and extrapolators are at 14%. The rest—everyone else—is

19% of investors. Because passive investors and risk avoiders have symmetric responses to the two different questions about S&P 500 moves (as does the "everyone else" category to an extent), the group of individual investors, taken as a whole, appears to exhibit the extrapolative behavior that we discussed when comparing this group to financial advisors and institutions.

In Table 2.4 we look at demographic patterns across different clusters. We assume that the clusters are independent collections of investors and compare their demographic averages; tests for statistical significance are given in Table A.6 in the Appendix. There are no significant differences in gender. Passive investors are the oldest and risk avoiders are the youngest among the four clusters (statistically significant at the 0.1% level). Unsurprisingly, passive investors are also the ones that are most likely to be retired (significant at the 0.1% level). This is consistent with the results from Figure 2-7, which compared the degree of activity between generations. Here we have a more detailed breakdown, showing that the more active group, made up of younger investors, tends to exhibit risk avoidance behavior following large S&P 500 changes.

There are no significant differences in net worth across the clusters—except that the "everyone else" group has the highest average net worth (significant at the 5% level). An interesting distinction arises when we look at the proportion of investors in different clusters who use a financial advisor. Passive investors and risk avoiders are much less likely than extrapolators and "everyone else" to use an advisor (significant at the 1% level). This makes intuitive sense, given that a financial advisor would likely try to discourage a risk avoider from significantly decreasing equity allocation in response to all large changes, as well as likely encouraging some fraction of passive investors into more active allocation. Interestingly enough, the extrapolator cluster is the one most likely to use a financial advisor—even though, as mentioned earlier, advisors usually provide contrarian advice to these individuals. We do not have sufficient evidence to explain this finding. However, it may be that advisors often deal with clients who extrapolate, and consequently advise an "opposite" allocation strategy to mitigate the bias of these investors.

Another dimension on which we compare investors is the degree of satisfaction with their 2014 investment returns. Not surprisingly, risk avoiders were much less satisfied in comparison to everyone else (significant at the 0.1% level). This is likely explained by their risk avoidance: they saw large positive returns on their investments, and thus decreased their equity allocation—and missed out on very good market returns over the full year, the S&P 500 returning 13.5% in 2014. We note that financial advisors would be particularly beneficial to risk avoiders, since they would encourage investors to stay invested in the market despite large swings, and in this way earn higher returns over longer time horizons.

Finally, we compare risk aversion across the clusters. Passive investors are the most risk-averse cluster (significant at the 0.1% level when comparing both direct gamble preferences and risk aversion coefficients). This is most likely due to its composition of predominantly older investors, who are more risk-averse than younger cohorts, as seen earlier. Ironically, risk avoiders appear to be the least risk-averse cluster, although the differences are not significant compared to the "everyone else" group.

We briefly comment on clustering results for the financial advisor and institutional investor responses. Because the sample sizes for these groups are much smaller than for the individual investor sample (2,342 for advisors and 660 for institutions), we cannot perform the in-depth analysis that we do for individuals. However, we can still look at the largest cluster in each group. In particular, we use the responses from the two questions about hypothetical S&P 500 moves to form three clusters (data exploration found three to be the most useful number). Tables A.7 and A.8 in the Appendix outline the relative sizes of the clusters and the average responses within each cluster. We see that the largest cluster in each sample clearly corresponds to a contrarian reaction; it makes up 47% of the sample for advisors and 68% for institutions. This further confirms our findings from the previous section that advisors and institutions are predominantly contrarian.

2.5 Individual Investor Predictive Analytics

We now look at the problem of predicting asset allocation behavior. Since we only have survey data available, the best we can do is use the responses to one set of questions for predicting the responses to other questions. We are particularly interested in predicting the three questions related to allocation decisions: the timing of equity allocation decrease around the financial crisis, and the two questions about reactions to S&P 500 moves. Our conclusion is that for individual investors we able to marginally beat the baseline for the crisis allocation timing question and are able to get a significant increase (of 15%) over the baseline when predicting how people react to S&P 500 changes.

We have decided to focus only on the individual investor dataset, since its sample size of 7,000 makes it large enough to perform predictive analytics. While we have also considered prediction for financial advisors and institutional investors, we decided to not proceed with it due to the small sample size (2,342 and 660, respectively) and more importantly, the lack of features that could be used for prediction. Thus, throughout this whole section we will only be dealing with individual investors.

Table 2.3 lists the possible responses and their coding for the three questions we focus on. We use six features for prediction:

- Gender
- Generation
- Net worth
- Use of financial advisor
- Gamble preferences
- Region

The possible values for the first three features above are shown in Table 2.1. The *use of financial advisor* variable is an indicator specifying if an investor uses a financial advisor or not. The *gamble preferences* variable is the answer to the question on which one of six gambles in Table 2.2 is preferred. Finally, *region* is one of the seven possible regions the investor is located in, shown in Table A.2 in the Appendix.

For questions about the reaction to the S&P 500 moves, we use two more features: the response to the question about asset allocation around the financial crisis, as well as an indicator for if an investor was satisfied with their 2014 investment returns or not. These features cannot be used for the crisis allocation timing question, since otherwise a look-ahead bias would arise.

We randomly split the sample into 70% for training and 30% for testing. We consider three different models: multinomial regression, CART, and random forest. For differentiating which model performs better we use a standard accuracy measure, whereby we measure the percent of responses on the test set that we predict correctly.

For each of the three allocation cases we look at by how much the models beat the baseline prediction. We use a simple baseline where we predict the same response for each investor; this response is chosen as the most common one among the ones in the training set. In all three cases we find this most common response is the passive one, whereby the investor chooses to not change their equity allocation.

Table 2.5 lists prediction accuracy for the model. The baseline model has relatively high accuracy. In the case of the financial crisis allocation question, with six possible responses, the baseline has 44% accuracy. For the questions about reactions to S&P 500 moves, the baseline has 34% and 33% accuracy, which is also quite high, considering five different responses are possible. We also see that across all questions random forest performs best, yielding close to 50% prediction accuracy.

We next investigate the prediction accuracy for each question individually. Multinomial regression gives accuracy of 44%, in comparison to 49% for both CART and random forest. This is a marginal improvement over the baseline – however it is important to keep in mind that the baseline prediction is already quite high. Table 2.6 shows the confusion matrix for the random forest model predictions; we see that when we predict a passive response, we obtain good prediction accuracy. When we predict a decrease in equity allocation between the first half of 2008 and the first half of 2009, prediction accuracy is not as good.

For the two questions about the reaction to S&P 500 changes results are better. For both questions multinomial regression and CART provide accuracy between 42% and 44% on the test set. Random forest is more accurate, getting 49% of predictions right for the reaction to a fall in the S&P 500, and 48% for the reaction to an S&P 500 rise. This is a substantial increase of 15% over the baseline. The confusion matrices for random forest predictions, shown in Table 2.6, also look good. For each of the predicted responses and both questions, accuracy is above 40%, whcih means that predictions are consistently accurate across all predicted responses. The accuracy of the “do nothing" prediction is above 50%, while we also have high accuracy in predicting an increase in equity allocation (although the sample size corresponding to predictions for this latter response is small).

We look into how individual variables contribute to prediction accuracy. While multinomial regression would allow for better interpretability, it performs worse than other methods. Hence we again focus on the random forest model applied to the two questions

Table 2.5: Accuracy of Investor Asset Allocation Predictions

		Prediction Method			
		Baseline	Multinomial Regression	CART	Random Forest
Scenario	Financial Crisis	44%	46%	49%	49%
	S&P 500 Fall	34%	42%	44%	49%
	S&P 500 Rise	33%	44%	43%	48%

This table shows how well we can predict asset allocation decisions under different investment scenarios. *Baseline* is when the same outcome is predicted in all cases. This is compared to three models: multinomial regression, CART, and random forest. Each model is estimated on the training set and then applied to the testing set. We measure accuracy as percent of observations in the test set that we predict correctly. Three scenarios are considered. *Financial Crisis* is the response to the question about when (if at all) an investor decreased their equity allocation around the financial crisis. *S&P 500 Fall* and *S&P 500 Rise* correspond to the two questions about how an investor would change their equity allocation following large moves in the S&P 500. For each case, the model with the highest accuracy is highlighted.

about reactions to S&P 500 moves. For each of the eight variables considered, we measure the decrease in accuracy on the test set if that variable is excluded from the model in comparison to the case when all variables are included. Table 2.7 shows the results. We see that accuracy does not substantially decline if a particular variable excluded. The effect of the crisis allocation timing feature seems to be more substantial, with accuracy decreasing the most when this variable is excluded. This could suggest that this variable is particularly important for prediction, which makes sense because it has explicit information about past investor allocation decisions.

We conclude that we can predict expected equity allocation behavior reasonably well by using investor demographics and past portfolio information. Using the random forest model we get close to 50% accuracy in predicting the responses to questions on asset allocation under various market scenarios. In the case of reactions to large S&P 500 changes, this is a big improvement over the baseline (which gives just 33%-34% accuracy). It is important to note that only a limited number of features are available for prediction, and if we had more details about past investor trading history, our predictions would likely improve.

2.6 Conclusion

Using a comprehensive global survey, we have identified differences in the investment behavior and risk tolerance preferences of individual investors, financial advisors, and

Table 2.6: Confusion Matrices for Random Forest Predictions

Timing of Equity Allocation Decrease around Financial Crisis		Predicted					
		1	2	3	4	5	6
Actual	Second half of 2007 (1)	0	6	25	5	0	131
	First half of 2008 (2)	0	30	65	11	0	190
	Second half of 2008 (3)	0	26	102	9	1	212
	First half of 2009 (4)	0	5	51	23	0	121
	Second half of 2009 (5)	0	2	31	7	3	116
	No Significant Change (6)	0	14	43	7	0	864
	Accuracy	N/A	36%	32%	11%	75%	53%

Equity Allocation Change Following S&P 500 Fall		Predicted				
		-2	-1	0	1	2
Actual	Large Decrease (-2)	144	116	152	14	2
	Small Decrease (-1)	68	298	207	9	1
	Do Nothing (0)	43	113	537	11	3
	Small Increase (1)	50	80	138	33	3
	Large Increase (2)	13	16	18	12	19
	Accuracy	45%	48%	51%	42%	68%

Equity Allocation Change Following S&P 500 Rise		Predicted				
		-2	-1	0	1	2
Actual	Large Decrease (-2)	104	7	103	153	0
	Small Decrease (-1)	52	22	125	110	1
	Do Nothing (0)	33	8	490	141	0
	Small Increase (1)	36	8	182	387	2
	Large Increase (2)	10	1	41	76	9
	Accuracy	44%	48%	52%	45%	75%

This table presents confusion matrices for predicting responses to the three allocation questions discussed earlier. In each case a random forest model is used for prediction. Columns correspond to different *predicted* responses, whereas rows correspond to *actual* responses. For each predicted response, we list the accuracy of prediction, measured as the empirical probability that the predicted response matches the actual one.

Table 2.7: Random Forest Predictions: Excluding One Variable

Predicting Change to:		S&P 500 Fall	S&P 500 Rise
All Variables Included		49.1%	48.2%
Variable Excluded	Gender	48.5%	47.8%
	Generation	48.0%	47.2%
	Net Worth	48.8%	47.5%
	Use of Financial Advisor	46.4%	46.4%
	Satisfied with 2014 Returns?	48.0%	47.4%
	Gamble Preferences	47.4%	46.2%
	Region	47.6%	45.4%
	Crisis Allocation Timing	45.8%	44.6%

This table presents how much accuracy in prediction is lost if one variable is excluded from the random forest model. For each variable, we train the model using the remaining seven variables on the training set and then measure prediction accuracy on the test set. We perform this exercise for the two questions about reactions to S&P 500 moves. For each question, we highlight the cells corresponding to the two variables for which accuracy decreases by the largest amount.

institutional investors. Advisors and institutions exhibit contrarian strategies in their behavior whereby they tend to change equity allocation in the opposite direction of recent returns on the S&P 500. This reaction is particularly strong for institutional respondents, 71% of whom would increase equity allocation following a fall in the S&P 500, and 66% would decrease allocation following a rise. This behavior is not as pronounced among financial advisors, because a large proportion of advisors tend to act passively and not change allocation at all.

By asking for preferences among six hypothetical gambles we are able to estimate the risk aversion coefficient for each of the three groups. Consistent with general intuition, individual investors are by far the most risk-averse, whereas institutional investors are the least risk-averse. We also compare risk aversion across different individual demographic categories, and find that risk aversion increases with age and decreases with net worth.

We observe significant heterogeneity among individual investors in terms of their allocation decisions. Using a clustering algorithm we classify investors into four distinct types. The first corresponds to *passive investors* and makes up about 50% of the sample. The other three types are *risk avoiders*, *extrapolators*, and *everyone else* and are of similar size, ranging between 14% and 19% of all respondents. Risk avoiders tend to significantly decrease equity allocation following large changes in the S&P 500 (both positive and negative), while extrapolators shift allocation in the same direction as those changes.

Passive investors are, on average, older and more risk-averse than the other three types, while risk avoiders tend to consist of younger investors.

We further investigate the individual investor dataset by addressing the problem of classifying investor responses to the allocation decision questions based on demographics information. We find that employing random forests provides a classification accuracy of around 50%, an improvement of over 15% over the baseline for both allocation questions.

Our results using this novel survey dataset have important implications for future research. First, this data gives us new insight into the allocation decisions of market participants over medium-term time horizons. While there have been been extensive studies of short-term trading by individual and institutional investors (e.g. by Grinblatt and Keloharju (2000) and Griffin, Harris, and Topaloglu (2003)), few have looked at the broader decisions of asset allocation. The fact that institutions and advisors are largely contrarian in their allocation strategies, while investors are on average slightly extrapolative, may be important in understanding the process of strategic asset allocation and the trading between these different market participants on the asset class level.

Our other important insight comes from our breakdown of individual investors into clusters of different behavioral types. Our "risk avoider" type has seldom been documented in the literature, and it may prove to be a useful component of future economic models. Recent papers already incorporate extrapolators (e.g. Barberis et. al (2015a, 2015b)) into their models, as well as fundamental investors, which correspond to our institutional investors. It would be interesting to see the equilibrium dynamics of all four different individual investor types play out, as well as the dynamics of potentially contrarian institutions taking the other side of trades made by extrapolators.

Finally, it is of significant benefit to study the behavior of financial advisors, and in particular, why so many of them advocate contrarian strategies to their clients. It also appears that individuals who use a financial advisor are more likely to exhibit extrapolative behavior. Do they ignore financial advice, or are advisors trying to "balance out" their extrapolation with contrarian suggestions? Further surveys and studies of historical financial advice and associated client action are needed to answer these questions.

2.7 Appendix

In this appendix, we include the specific survey questions used for individual and institutional investors (Section 2.7.1), and for financial advisors (Section 2.7.2), and summarize the characteristics of the survey participants (Section 2.7.3).

2.7.1 Individual Investor and Institutional Survey Questions

The behavioral questions presented in the Individual Investor and the Institutional Surveys were exactly the same. They are listed below.

1. Of the following six gambles, which would you prefer the most?
   a. Win $28,000 with probability 100%
   b. Win $36,000 with probability 50%; win $24,000 with probability 50%
   c. Win $44,000 with probability 50%; win $20,000 with probability 50%
   d. Win $52,000 with probability 50%; win $16,000 with probability 50%
   e. Win $60,000 with probability 50%; win $12,000 with probability 50%
   f. Win $70,000 with probability 50%; win $2,000 with probability 50%

2. How would you change your asset allocation if the S&P 500 declined between 10% and 20% during the next six months and other asset classes performed as you expected?
   a. I would do nothing
   b. I would decrease my stock or shares allocation slightly
   c. I would decrease my stock or shares allocation significantly
   d. I would increase my stock or shares allocation slightly
   e. I would increase my stock or shares allocation significantly

3. How would you change your asset allocation if the S&P 500 declined between 10% and 20% during the next six months and other asset classes performed as you expected?
   a. I would do nothing
   b. I would decrease my stock or shares allocation slightly
   c. I would decrease my stock or shares allocation significantly
   d. I would increase my stock or shares allocation slightly
   e. I would increase my stock or shares allocation significantly

4. Around what time during the Financial Crisis of 2007–2009 did you significantly decrease your allocation to stocks in your investment portfolio?
   a. Second half of 2007
   b. First half of 2008
   c. Second half of 2008
   d. First half of 2009
   e. Second half of 2009
   f. I did not significantly decrease my allocation to stocks during the Financial Crisis

2.7.2 Financial Advisor Survey Questions

The behavioral questions presented in the Financial Advisor Survey were essentially the same as the questions for individual investors. For questions pertaining to asset allocation, advisors were asked about how their clients changed their allocations/how they would advise clients to change their allocations rather than how advisors would alter allocation in their personal investment portfolios. We also asked a question about when clients typically increased equity allocation after the Financial Crisis. Since this question was not asked to the individual investors or institutional participants, we did not include its results in our analysis.

1. Of the following six gambles, which would you prefer the most?
   a. Win $28,000 with probability 100%
   b. Win $36,000 with probability 50%; win $24,000 with probability 50%
   c. Win $44,000 with probability 50%; win $20,000 with probability 50%
   d. Win $52,000 with probability 50%; win $16,000 with probability 50%
   e. Win $60,000 with probability 50%; win $12,000 with probability 50%
   f. Win $70,000 with probability 50%; win $2,000 with probability 50%

2. Around what time after the Financial Crisis of 2007–2009 did most of your clients significantly decrease their allocation to stocks in their investment portfolios?
   a. Second half of 2008
   b. First half of 2009
   c. Second half of 2009
   d. First half of 2010
   e. Second half of 2010
   f. Most of my clients did not significantly increase their allocation to stocks after the Financial Crisis.

3. How would you advise your clients to change their asset allocation if the S&P 500 declined between -10% and -20% during the next six months and other asset classes performed as expected?
   a. I would advise clients to do nothing
   b. I would advise clients to decrease equity allocation slightly
   c. I would advise clients to decrease equity allocation significantly
   d. I would advise clients to increase equity allocation slightly
   e. I would advise clients to increase equity allocation significantly

4. How would you advise your clients to change their asset allocation if the S&P 500 increased between -10% and -20% during the next six months and other asset classes performed as expected?
   a. I would advise clients to do nothing
   b. I would advise clients to decrease equity allocation slightly
   c. I would advise clients to decrease equity allocation significantly

d. I would advise clients to increase equity allocation slightly

e. I would advise clients to increase equity allocation significantly

5. Around what time during the Financial Crisis of 2007–2009 did most of your clients significantly decrease their allocation to stocks in their investment portfolios?

a. Second half of 2007

b. First half of 2008

c. Second half of 2008

d. First half of 2009

e. Second half of 2009

f. Most of my clients did not significantly decrease their allocation to stocks during the Financial Crisis.

2.7.3 Survey Respondents Characteristics

We provide more details on the subjects included in the surveys. Table A.1 and A.2 show the country and region breakdowns of respondents for the Individual Investor Survey, while Table A.3 does this for the Financial Advisor Survey. In Table A.4 we list the types of institutions included in the Institutional Investor survey. Finally, Table A.5 presents the definitions of the individual investor demographic categories.

Country	**# Respondents**	**Country**	**# Respondents**
Argentina	200	Japan	350
Australia	250	Mexico	350
Canada	250	Singapore	500
Chile	200	Spain	500
Colombia	200	Switzerland	350
France	500	UAE/Qatar/Kuwait	350
Germany	500	United Kingdom	750
Hong Kong	500	United States	750
Italy	500		

Table A.1: Number of respondents, by country, in the Individual Investor Survey.

Region	**# Respondents**	**Region**	**# Respondents**
Asia & The Pacific	1600	North America	250
Europe	2350	UK	750
Latin America	950	US	750
Middle East	350		

Table A.2: Number of respondents, by region, in the Individual Investor Survey.

Country	# Respondents	Country	# Respondents
Canada	150	Singapore	146
Chile	137	Spain	150
France	150	Switzerland	150
Germany	150	UAE/Qatar/Kuwait	144
Hong Kong	144	United Kingdom	300
Italy	150	Uruguay	140
Panama	131	United States	300

Table A.3: Number of respondents, by country, in the Financial Advisor Survey.

Institution Type	# Respondents
Central Bank	11
Corporate Pension Plan	196
Insurance Company	100
Non-Profit (Endowment/Foundation)	131
Public/Government Pension Plan	140
Sovereign Wealth Fund	69
Other Institution	13

Table A.4: Breakdown of respondents by institution type, in the Institutional Investor Survey.

Demographic	Definition
Generation Y	18–33 years old
Generation X	34–49 years old
Baby Boomers	50–68 years old
Pre-Baby Boomers	69 years old and above
Mass Market	NW: $200,000 - $300,000
Mass Affluent	NW: $300,000 - $500,000
Emerging HNW	NW: $500,000 - $1,000,000
High Net Worth	NW: $1,000,000 and above

Table A.5: Descriptions of the different demographic categories used in the Individual Investor Survey. The abbreviation *NW* means Net Worth, while *HNW* means High Net Worth.

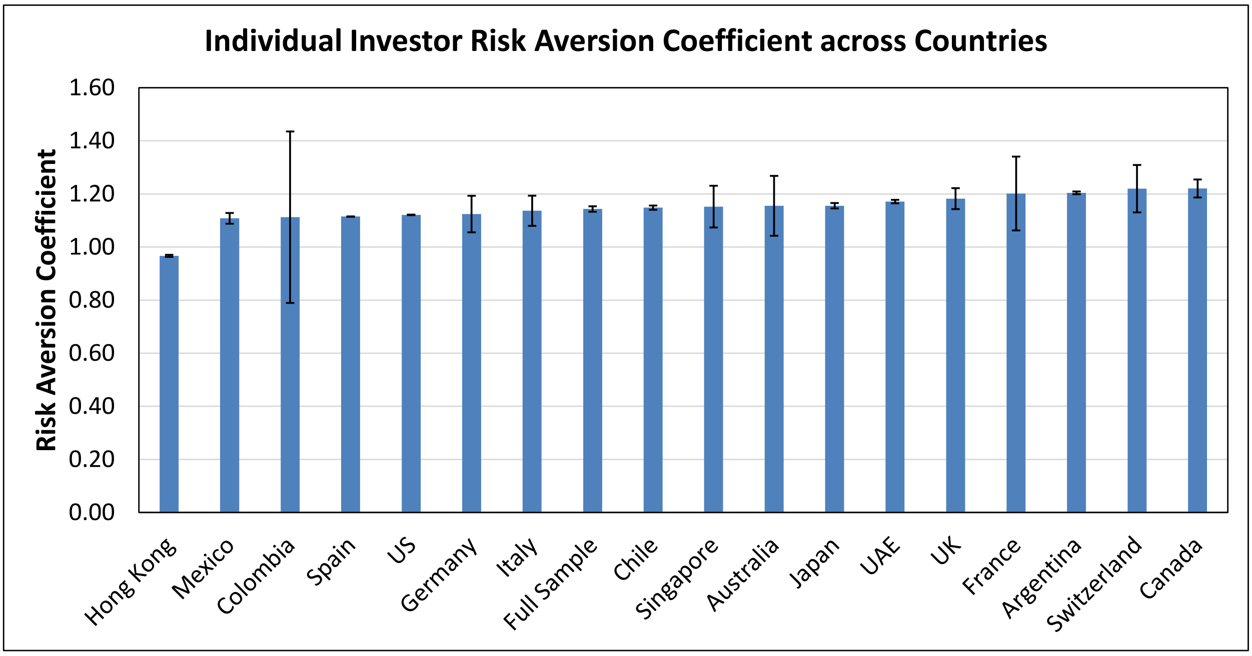

Figure 2.7.1: Estimated risk aversion coefficients for individual investors across countries. Error bars correspond to one standard error. Standard errors are large for some countries due to small sample size, e.g. for Colombia we have only 200 observations.

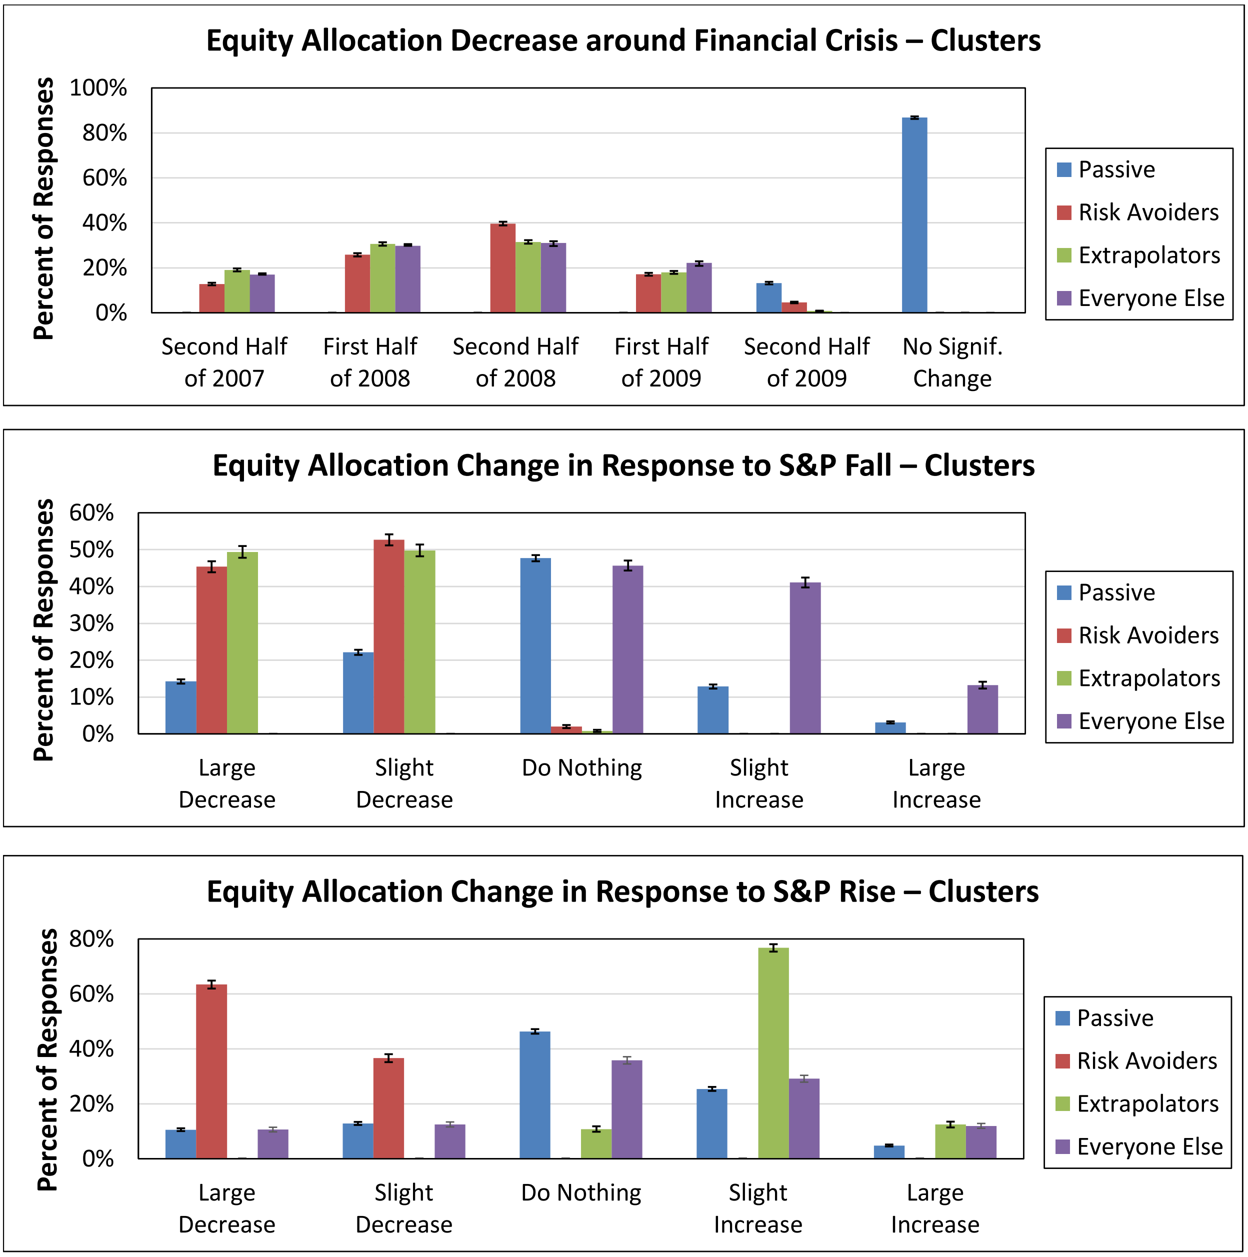

Figure 2.7.2: Distributions of responses for the four clusters created from the Individual Investor Survey. Error bars correspond to one standard error.

Statistical Significance for Comparisons of Mean Demographics across Clusters

Gender	Passive	Risk Av.	Extrap.	Other
Passive	0.500	0.245	0.496	0.456
Risk Av.	0.755	0.500	0.703	0.691
Extrap.	0.504	0.297	0.500	0.470
Other	0.544	0.309	0.530	0.500

Satisfied 2014	Passive	Risk Av.	Extrap.	Other
Passive	0.500	1.000	0.974	0.064
Risk Av.	0.000	0.500	0.000	0.000
Extrap.	0.026	1.000	0.500	0.002
Other	0.936	1.000	0.998	0.500

Age	Passive	Risk Av.	Extrap.	Other
Passive	0.500	1.000	1.000	1.000
Risk Av.	0.000	0.500	0.000	0.000
Extrap.	0.000	1.000	0.500	0.977
Other	0.000	1.000	0.023	0.500

Retired	Passive	Risk Av.	Extrap.	Other
Passive	0.500	1.000	1.000	1.000
Risk Av.	0.000	0.500	0.773	0.026
Extrap.	0.000	0.227	0.500	0.004
Other	0.000	0.974	0.996	0.500

Net Worth	Passive	Risk Av.	Extrap.	Other
Passive	0.500	0.635	0.615	0.010
Risk Av.	0.365	0.500	0.488	0.017
Extrap.	0.385	0.512	0.500	0.022
Other	0.990	0.983	0.978	0.500

Gamble Pref.	Passive	Risk Av.	Extrap.	Other
Passive	0.500	0.000	0.000	0.000
Risk Av.	1.000	0.500	0.998	0.626
Extrap.	1.000	0.002	0.500	0.005
Other	1.000	0.374	0.995	0.500

Advised	Passive	Risk Av.	Extrap.	Other
Passive	0.500	0.188	0.000	0.000
Risk Av.	0.812	0.500	0.000	0.003
Extrap.	1.000	1.000	0.500	0.986
Other	1.000	0.997	0.014	0.500

Risk Aversion	Passive	Risk Av.	Extrap.	Other
Passive	0.500	0.000	0.000	0.000
Risk Av.	1.000	0.500	1.000	0.863
Extrap.	1.000	0.000	0.500	0.005
Other	1.000	0.137	0.995	0.500

Table A.6: Hypothesis tests for statistical significance when comparing means of demographic categories across the four clusters from the Individual Investor Survey. We assume that each cluster corresponds to an independent collection of investors. For every demographic, we calculate the mean response within each cluster. The resulting means are shown in Table 2.4.
For every two means, we test for statistical significance using Welch's t-test—with the exception of the risk aversion coefficient, where we employ a z-test using the risk aversion coefficients and associated standard errors from the estimation. Each cell contains the p-value associated with testing if the column cluster mean minus the row cluster mean is greater than zero. A cell is colored green if the mean in the row cluster is significantly less than the mean in the column cluster, at the 5% significance level. A cell is colored red if the mean in the row cluster is significantly greater than the mean in the column cluster, at the 5% significance level.
Risk Av. means *Risk Avoiders* cluster; *Extrap.* is the *Extrapolators* cluster. *Satisfied 2014* is an indicator for if an investor was satisfied with their 2014 investment returns. *Risk Aversion* is short for the risk aversion coefficient.

Clustering of Financial Advisors			
	Cluster 1	**Cluster 2**	**Cluster 3**
% Respondents	47%	37%	16%
	Allocation Decisions		
Reaction to S&P Fall	1.1	0.4	-1.4
Reaction to S&P Rise	-1.1	0.2	0.3

Table A.7: Clustering of allocation decision responses from the Financial Advisor Survey. For each cluster, we present the percent of respondents and mean response based on the response coding in Table 2.3.

Clustering of Institutional Investors			
	Cluster 1	**Cluster 2**	**Cluster 3**
% Respondents	68%	22%	10%
	Allocation Decisions		
Reaction to S&P Fall	1.5	-0.6	-1.4
Reaction to S&P Rise	-1.1	0.7	-1.7

Table A.8: Clustering of allocation decision responses from the Institutional Investor Survey.

Chapter 3

Algorithmic Models of Investor Behavior

(joint work with Andrew W. Lo)

Abstract

We propose a heuristic approach to modeling investor behavior by simulating combinations of simpler systematic investment strategies associated with well-known behavioral biases—with functional forms motivated by an extensive review of the behavioral finance literature review—with parameters calibrated from historical data. We compute the investment performance of these heuristics individually and in pairwise combinations using both simulated and historical asset-class returns. The mean-reversion or momentum nature of a heuristic can often explain its effect on performance, depending on whether asset returns are consistent with such dynamics. We propose a Markov chain Monte Carlo (MCMC) algorithm for estimating the parameters of these strategies and their implicit weights, and show that this method can successfully infer the relative importance of each heuristic among a large cross-section of investors, even when the number of observations per investor is quite small. We also compare the accuracy of the MCMC approach to regression analysis in predicting the relative importance of heuristics at the individual and aggregate level, and conclude that MCMC predicts aggregate weights more accurately while regression outperforms in predicting individual weights.

Keywords: Portfolio Management; Asset Allocation; Performance Attribution; Behavioral Finance.

JEL Classification: G02, G11, G12

3.1 Introduction

Individuals often behave in predictable ways when making investment decisions, in patterns that are sometimes called "behavioral biases". These behavioral biases can be quite costly to investors, leading to large losses and inappropriate amounts of risk in their portfolios. However, in other circumstances, these same biases can lead to positive outcomes, such as protecting investors from even larger losses than they would otherwise experience. A clearer understanding of these biases would benefit all investors, allowing them to modify their behavior according to their changing circumstances.

Over the past 40 years, academics have made considerable progress in empirically documenting these biases in financial decision-making, both by running experiments in a laboratory setting and by analyzing historical data on investor trades. However, this work has mainly focused on one or two biases in isolation, and only a few studies have considered the full spectrum of their impact on realized investment performance. Furthermore, while there is strong evidence that specific patterns of behavior exist, the problem of their origin and evolution over time remains open.

On the theoretical side, academics have proposed several behavioral models to complement the original paradigm of the rational investor. These models are designed to capture particular irrational aspects of financial behavior. They can be used to generate optimal trading strategies for investors influenced by these behaviors, and they imply potential changes to the financial markets under their influence. Prospect theory, put forward by Kahneman and Tversky (1992), is perhaps the most popular model of this type. However, even this simple framework is challenging to apply to the prediction of investor behavior since there are several free parameters that must be calibrated, e.g., the reference point and the investment horizon.

We propose a novel framework to analyze and predict investor behavior in which behavior is modeled as a combination of simpler heuristics, each linked to one or more well-known biases such as loss aversion or overconfidence and motivated by the existing academic and practitioner literature. For tractability and ease of interpretation, we use very few parameters to describe each heuristic, and these heuristics imply specific strategies that, when systematically combined, form a portfolio that can describe the trading behavior of an investor. This approach is purely algorithmic and can be viewed as an "artificial intelligence" model of investor behavior. It can be applied to individual trading data from which an individual's heuristics can be calibrated, thereby approximating the individual's behavior algorithmically. This framework can also be used to analyze investor behavior under various conditions, and can predict future investor behavior.

We begin with explicit definitions of our heuristics, all of them motivated by existing examples of behavioral biases in the literature. After specifying our investor decision-making framework, we investigate how pairs of biases affect investment performance under different specifications for the risky asset's return-generating process. We also investigate the interactions among these heuristics when all of them are active. We perform an analogous exercise with historical data, applying various heuristics to different asset classes, including equities, bonds, foreign exchange, and commodities. Despite the relative simplicity of our approach, we are able to capture important aspects of investor behavior and corresponding biases.

We then show how to estimate the parameters of the heuristics from trading data by applying Markov chain Monte Carlo (MCMC) methods to cross-sectional data on investment decisions data, assuming several stylized cases: when only one heuristic is active, when a pair of heuristics is active, and when all heuristics are active. With only 100 observations for each investor, and with a relatively small number of investors overall (1,000), we are able to estimate several important parameters with reasonable accuracy: the default strategy, the allocation noise, and the heuristic weights. It is important to note that we have less success at estimating time horizons and the thresholds for the heuristics.

Finally, we study the problem of predicting asset allocation decisions within our framework. We simulate data according to our heuristics, estimate the model parameters on a training subset of the simulated data, and then attempt to predict subsequent allocation decisions using these estimates on an out-of-sample basis. This approach is able to successfully predict aggregate investor behavior, and it usually outperforms regression-based predictions. However, at the level of the individual investor, regressions tend to have more accurate out-of-sample predictions.

3.2 Decision-Making Framework

In our model, an investor follows several decision-making rules or heuristics. These rules are motivated by existing evidence on behavioral biases in investor behavior. While some of these rules may serve as proxies for particular investor biases, it is important to remember that they do not fully capture them. Rather, these rules can be thought of as generic versions of common decision-making patterns, which have strong connections with particular behavioral biases in the literature.

Throughout the rest of the paper, we will use the words “rule" and “heuristic" interchangeably.

We consider an agent investing over time in risky assets $1, 2, \ldots, N$ and a risk-free asset $N+1$. For tractability we will initially assume $N = 1$, so that there is only one risky asset. The agent follows a default strategy which invests a portion $x_{i,t}$ of the portfolio into asset i at the start of period t.

Besides the default strategy, the investor is subject to heuristics $1, 2, \ldots, K$ that result in adjustments to the portfolio. We view each heuristic j as a separate (systematic) investment strategy which specifies the portfolio weight z_t^j to be invested into the risky asset at the start of period t. The strategies are combined as a weighted average, with weight $\gamma_{j,t}$ assigned to heuristic j. Thus the portfolio allocation to the risky asset at the start of period t is:

$$w_{1,t} = x_{1,t} + \sum_{j=1}^{K} \gamma_{j,t} z_t^j$$

The allocation to the risk-free asset is just $w_{2,t} = 1 - w_{1,t}$.

For each heuristic, the decision on the portfolio adjustment z_j^t is made at the end of period $t-1$ using the asset prices and portfolio performance up to date $t-1$, inclusive. This decision consists of two parts. The first is the “heuristic trigger" – the conditions that need to occur to trigger the heuristic. If a heuristic is triggered, it is assigned a strength of s_t^j; this strength is set to 0 when the heuristic is inactive. The second component is the “heuristic trade" – the proposed asset allocation z_t^j as a result of the heuristic. Each of these two time-varying variables are functions of past performance and asset returns with fixed parameters.

We will use the following notation:

- $P_{i,t}$ – price of asset i at the end of period t
- $B_{i,t}$ – amount of portfolio held in asset i at the end of period t
- B_t – total portfolio balance at the end of period t

The following parameters will be used for describing the biases (different heuristics will have different parameters):

- L – horizon over which past performance is measured
- α – gain/loss threshold
- h – horizon of the heuristic
- n_o – number of remaining periods of the heuristic

In the next few sections, we summarize the relevant behavioral finance literature used to motivate our heuristics, and present our heuristic definitions.

3.2.1 The Disposition Effect

People tend to hold on to losers for too long and to sell winners too soon, a behavior known as the *disposition effect*, first discussed by Shrefin and Statman (1985). Odean (1998) documents the disposition effect more explicitly by examining trading records of individual investors. Further evidence of this effect in both an experimental setting and using trading records has been produced by Weber and Camerer (1998), Garvey and Murphy (2004), and Frazinni (2006).

Several explanations of the disposition effect have been put forward. Fogel and Berry (2006) propose that regret is the primary explanation for why people hold onto losing investments and sell winning ones too soon. Barberis and Xiong (2009, 2012) show that prospect theory alone is not always able to explain the disposition effect, and propose an alternative model where investor utility depends on *realized* gains. This would cause an investor to be unwilling to sell a losing investment due to its associated realized loss, and instead would wait until it becomes profitable.

Another potential cause of this effect is self-control, or rather, its lack. Shefrin and Statman (1985) suggest that self-control could be used to explain the disposition effect.

People who lack self-discipline and prefer short-term gratification may prematurely sell off winners, while investors may hold on to losers for too long, since they cannot exercise enough discipline to sell them. This may be why some traders use set trading rules, such as stop-loss orders, in order to prevent such mistakes. Traders are often advised to exercise good self-control in order to ride out their profits, in spite of the temptation to sell too soon; see, for example, Krivo (2012).

Finally, loss aversion also contributes to the disposition effect; we will discuss this in the next section.

There are two basic behaviors described by the disposition effect. The first is selling profitable investments too early. Here, soon after initiating the position, an investor sees a stock perform well and sells it, instead of holding it for potentially larger gains. We model a *short-term gains* heuristic by setting the behavior to be triggered if the gains in the investor portfolio over a particular (short) time horizon exceed a prespecified threshold.

Definition 3. *A* ***short-term gains*** *heuristic is triggered as follows:*

$$(s_t, n_o) = \begin{cases} (1, h) \text{ if } \dfrac{B_t}{B_{t-L}} > \alpha \\ (1, n_o - 1) \text{ otherwise and if } n_o > 1 \\ (0, 0) \text{ otherwise} \end{cases} \tag{1}$$

The heuristic trade is:

$$z_{t+1} = \begin{cases} -1 \text{ if } s_t = 1 \\ 0 \text{ if } s_t = 0 \end{cases} \tag{2}$$

The second aspect of the disposition effect is the reluctance to exit losing positions. For example, an investor buys a stock expecting its price to increase. Instead, its price falls. The investor hopes the negative trend will reverse, and keeps holding the stock, even though it can potentially further decline in value. In our model, a *short-term losses* heuristic is triggered if the losses over a particular (short) horizon exceed a prespecified threshold. In that case, the investor is inclined to increase his position.

Definition 4. *A **short-term losses** heuristic is triggered as follows:*

$$(s_t, n_o) \quad = \quad \begin{cases} (1, h) \text{ if } \dfrac{B_t}{B_{t-L}} < \alpha \\ (1, n_o - 1) \text{ otherwise and if } n_o > 1 \\ (0, 0) \text{ otherwise} \end{cases} \tag{3}$$

The heuristic trade is:

$$z_{t+1} \quad = \quad \begin{cases} 1 \text{ if } s_t = 1 \\ 0 \text{ if } s_t = 0 \end{cases} \tag{4}$$

3.2.2 Loss Aversion

Loss aversion is the tendency of investors to avoid losses. This behavioral anomaly affects people in several ways. First, they are reluctant to accept investment opportunities with a large downside relative to gains. Second, investors are reluctant to liquidate a position showing a paper loss, the disposition effect discussed in the previous section. Third, if investors have suffered a large loss, they become even more frightened of losing money, and often stay out of risky markets for a considerable period of time. In our model, we focus on describing this last phenomenon.

The literature on loss aversion is extensive. Kahneman and Tversky (1992) begin by conducting a classic series of experiments demonstrating loss aversion, offering subjects a series of gambles, and asking how much they would be willing to pay for each gamble. Kahneman and Tversky conclude that for the participants, losses matter about twice as much as gains of comparable magnitude. Further experimental evidence of this phenomenon has been produced by Thaler et al. (1997), demonstrating that people care more about losses than about gains, especially if their performance is evaluated frequently, and by Haigh and List (2005), who show that professional traders also exhibit loss aversion.

People tend to invest less in the markets after experiencing losses. This is probably due to a combination of loss aversion and other behaviors, such as the extrapolation of past returns or the negative emotions following a loss. Strahilevitz, Odean, and Barber (2011) analyze individual trading records and find that investors are less likely to repurchase stocks they had previously sold for a loss. Malmendier and Nagel (2009) show that investors who experienced poor returns on equities during their lifetime tend to invest less in stocks. Benartzi (2001) documents that employees allocate 30% less of their discretionary 401(k) contributions to company stock if the stock has performed poorly over the past 10 years, compared to the employees of firms whose stock has done well.

We model loss aversion as a *long-term losses* heuristic to be triggered if portfolio losses over a particular (long) horizon exceed a prespecified threshold. In that case, the investor will want to decrease his position in the security.

Definition 5. *A **long-term losses** heuristic is triggered as follows:*

$$(s_t, n_o) \quad = \quad \begin{cases} (1, h) \text{ if } \dfrac{B_t}{B_{t-L}} < \alpha \\ (1, n_o - 1) \text{ otherwise and if } n_o > 1 \\ (0, 0) \text{ otherwise} \end{cases} \tag{5}$$

The heuristic trade is:

$$z_{t+1} \quad = \quad \begin{cases} -1 \text{ if } s_t = 1 \\ 0 \text{ if } s_t = 0 \end{cases} \tag{6}$$

3.2.3 Overconfidence

People are overconfident. They tend to overestimate their ability to do certain tasks, [1] to think they are better than other people, [2] and to be unrealistically confident in the accuracy of their beliefs. [3] This behavior has been observed in many different areas, including entrepreneurship (Koellinger, Minniti, Schade, 2007), medicine (Berner and Graber, 2008), driving (Svenson, 1981), and negotiation (Neale and Bazerman, 1985). Not surprisingly, it also shows up in finance.

Overconfident traders often believe they can beat the market because they overestimate how much money they will make on their trades. This results in a higher trading volume, and usually in lower investor profits due to excessive trading, as documented by Odean (1998a), and by Barber and Odean (2001). Odean (1998a) concludes that investors do not merely overestimate the accuracy of their beliefs, but are also overconfident about their ability to interpret the information used to arrive at their beliefs.

A few empirical studies have been conducted to analyze how overconfidence arises

[1] Buehler, Griffin, and Ross(1994), Nowell and Alston (2007), Cassar (2010)
[2] Alicke et al. (1995), Camerer and Lovallo (1999), Guenther and Alicke (2010)
[3] Fischhoff et al. (1977), Klayman et al. (1999), Jørgensen et al. (2004)

from past returns and how this affects subsequent trading. Glaser and Weber (2009) study trading records of individual investors to find that, after experiencing high portfolio returns, investors tend to overweigh higher risk stocks, and reduce the number of stocks in their investment portfolio. Chuang and Susmel (2011) find that both individual and institutional investors trade more after good market returns in bull markets, in momentum markets, and in low volatility environments. Deaves, Luders, and Luo (2009) analyze students trading in a laboratory setting, and find that overconfidence due to overestimation of the accuracy of one's beliefs, as well as the belief that one is better than average, leads to more trading.

We focus on a phenomenon closely related to overconfidence, called self-attribution. When investors make a profitable trade, they tend to attribute their success to pure skill, and to underestimate the element of luck involved. As a result, they become more confident in their abilities following a winning streak of investments – and sometimes suffer large losses afterwards because they take excessively large positions. Two prominent papers analyzing this behavior within a theoretical framework are Daniel, Hirshleifer, and Subrahmanyam (1998) and Gervais and Odean (2001).

In our model a *long-term gains* heuristic is triggered if the investor's profits over a particular (long) horizon exceed a prespecified threshold, after which the investor increases his position.

Definition 6. *A* ***long-term gains*** *heuristic is triggered as follows:*

$$(s_t, n_o) = \begin{cases} (1, h) \text{ if } \dfrac{B_t}{B_{t-L}} > \alpha \\ (1, n_o - 1) \text{ otherwise and if } n_o > 1 \\ (0, 0) \text{ otherwise} \end{cases} \tag{7}$$

The heuristic trade is:

$$z_{t+1} = \begin{cases} 1 \text{ if } s_t = 1 \\ 0 \text{ if } s_t = 0 \end{cases} \tag{8}$$

3.2.4 The Gambler's and Hot Hand Fallacies

People incorrectly predict future outcomes after observing random sequences. In some situations, people believe in mean reversion (a heuristic known as the *gambler's fallacy*), while in others, they believe in momentum (the *hot hand fallacy*).

The gambler's fallacy tends to be associated with an incorrect interpretation of randomness in a random sequence. Rabin and Vayanos (2010) estimate that people observing flips of a fair coin on average estimate a 70% probability of a tail on the next flip after observing three consecutive heads. Oskrasson et al. (2009) provide a detailed summary of the existing literature on people's judgment of random sequences. It is consistently found that randomness tends to be associated with the belief that a trend in observed values will end. Empirical evidence supporting the gambler's fallacy has been documented among lottery players (Clotfelter and Cook, 1993), casino gamblers (Cronson and Sundali, 2005), and institutional investors (Shefrin, 2000). A belief in mean reversion may also partly explain the disposition effect.

The hot hand fallacy occurs when people observing a random sequence believe it is actually not random. Gilovich, Vallone, and Tversky (1985) find that people erroneously believe in "hot streaks" in basketball, while Cronson and Sundali (2005) show that some casino players increase bets after roulette winnings, thinking their streak will continue. Burns (2003) and Johnson, Tellis, and MacInnis (2005) document the hot hand fallacy among investors in laboratory experiments.

While both the gambler's fallacy and the hot hand fallacy are closely related to the heuristics discussed previously, they are distinct because they are related to sequences of binary events, rather than changes in prices or wealth. Thus, in our model we will look at the sign of returns over a specified number of periods L. If at least m of these returns are positive, or at least m are negative, then a *binary streak* heuristic is triggered.

We assume the heuristic lasts for only one period, so we have no need to keep track how long the heuristic will remain in effect. However, we record the direction d, which specifies whether most of the returns over the past horizon are positive or negative. We will not speculate which of the two heuristics will dominate (considering they work in different directions). Rather, we assume the investor follows the gambler's fallacy, and note that we still capture the hot hand fallacy if the weight assigned to the heuristic is negative.

Definition 7. *A **binary streak** heuristic is triggered as follows:*

$$(s_t, d) = \begin{cases} (1,1) \text{ if } \sum_{v=t-L+1}^{t} \mathbb{I}\{P_{1,v} > P_{1,v-1}\} \geq m \\ (1,-1) \text{ if } \sum_{v=t-L+1}^{t} \mathbb{I}\{P_{1,v} < P_{1,v-1}\} \geq m \\ (0,0) \text{ otherwise} \end{cases} \tag{9}$$

The heuristic trade is:

$$z_{t+1} = \begin{cases} -1 \text{ if } d = 1 \text{ and } s_t = 1 \\ 1 \text{ if } d = -1 \text{ and } s_t = 1 \\ 0 \text{ if } s_t = 0 \end{cases} \tag{10}$$

3.2.5 Regret

The regret heuristic is less basic than the heuristics described earlier, as it deals with the reaction of an investor observing the price behavior of a risky asset relative to his position. This heuristic captures the common idea that sometimes an investor “misses out" on potential gains by not investing in the asset when it went up in value. This investor will likely suffer a negative emotion, but his financial reaction is not obvious. On the one hand, he may feel disappointed in the stock, and no longer want to invest in it. On the other hand, he may wish to make up for his “mistake" by investing in the asset instead. Strahilevitz, Odean, and Barber (2011) analyze retail trading records and conclude that the former effect is more common, at least for the situation where an investor sells a stock that goes up in price following the sale. Weber and Welfens (2011) obtain similar results in a laboratory setting.

Lin, Huang, and Zeelenberg (2006) surveyed investors to ask which hypothetical situations they would feel the most regret as a result of their actions. They find that non-investment in a profitable stock has the largest impact on regret. Also, as mentioned above, the desire to avoid regret may be the primary cause of the disposition effect (see for example Fogel and Berry, 2006).

We will model regret by looking at situations where the return on a risky asset over a specified horizon is high (above a threshold α), while the return on the investor portfolio is low (below a threshold θ). This will cause an investor to experience regret, and to be more inclined to invest in the asset. Our specification of the *missing out* heuristic also incorporates the possibility of an investor decreasing his position in the asset instead – this will happen if the weight assigned to this heuristic is negative.

Definition 8. *A **missing out** heuristic is triggered as follows:*

$$(s_t, n_o) = \begin{cases} (1, h) \text{ if } \dfrac{B_t}{B_{t-L}} < \alpha \text{ and } \dfrac{P_{1,t}}{P_{1,t-L}} > \theta \\ (1, n_o - 1) \text{ otherwise and if } n_o > 1 \\ (0, 0) \text{ otherwise} \end{cases} \tag{11}$$

The heuristic trade is:

$$z_{t+1} = \begin{cases} 1 \text{ if } s_t = 1 \\ 0 \text{ if } s_t = 0 \end{cases} \tag{12}$$

3.2.6 Anchoring

The final heuristic we consider is *anchoring.* In this heuristic, investors evaluate information by comparing it to an important reference point. For example, a trader may decide on whether to buy or sell a stock based on its current price relative to its purchase price, a comparison which also may yield the disposition effect (Shapira and Venezia, 2001). While the purchase price is important, other aspects of past security and portfolio performance also may influence the reference point. Baucells, Weber, and Welfens (2011) estimate various functional forms for the reference point based on subject behavior in a laboratory setting. They propose a model which uses a linear combination of five prices: the current price, the purchase price, the average of prices over the holding period, as well as the high and the low prices. They find that the coefficients for all five variables are statistically significant in determining the anchor point.

George and Hwang (2004) suggest that the 52-week high stock price may be used by some investors as a natural anchor, explaining a large part of momentum strategy profits. More generally, investors may often compare the stock price to a recent high, and buy the stock because it "looks cheap", even if it is fairly priced based on its fundamentals. This is known among practitioners as a *value trap* (Yee, 2008, Speece and Rogers, 2010).

We will model this last aspect of anchoring by considering situations where an investor sees that the stock price has declined by more than a prespecified threshold relative to the high of the stock within a past time horizon. If that happens, under the *recent decline* heuristic, the investor is inclined to increase his position in the stock.

Definition 9. *A* ***recent decline*** *heuristic is triggered as follows:*

$$(s_t, n_o) \quad = \quad \begin{cases} (1, h) \text{ if } \dfrac{P_{1,t}}{max(P_{1,t-L+1}, P_{1,t-L+2}, \ldots, P_{1,t})} < \alpha \\ (1, n_o - 1) \text{ otherwise and if } n_o > 1 \\ (0, 0) \text{ otherwise} \end{cases} \tag{13}$$

The heuristic trade is:

$$z_{t+1} \quad = \quad \begin{cases} 1 \text{ if } s_t = 1 \\ 0 \text{ if } s_t = 0 \end{cases} \tag{14}$$

Having defined the component heuristics of an investor's decision-making process, we are ready to apply them to various investment scenarios. We start by outlining the parameters of our model, both for individual heuristics and for their weights in a portfolio. We then examine how investor decisions will affect this portfolio, simulating risky asset prices for different return-generating processes. After that, we consider the historical performance of investments in various asset classes while subject to these heuristics.

3.3 Model Parameters

Recall our specification for the investor's allocation problem:

$$w_{1,t} = x_{1,t} + \sum_{j=1}^{K} \gamma_{j,t} z_t^j$$

Each of the variables on the right side at time t is a function of the information $\mathcal{F}_{t-1}$ available to the investor after the end of period $t-1$; this includes past portfolio balances B_s and asset prices $P_{i,s}$ for $s \leq t-1$. These variables are functions with parameters that stay fixed throughout the whole investment horizon:

$$x_{1,t} = x_{1,t}(\alpha, \mathcal{F}_{t-1}), \;\; \gamma_{j,t} = \gamma_{j,t}(\beta_j, \mathcal{F}_{t-1}), \;\; z_t^j = z_t^j(\theta_j, \mathcal{F}_{t-1})$$

For simplicity, we will assume that the default strategy $x_{1,t}$ used by an investor maintains a constant allocation, which corresponds to the "buy and hold" strategy. In this case, the vector α contains just one parameter – which is the level, so that $x_{1,t} = \alpha$. Similarly, we will assume that the weights assigned to the heuristics are not time-varying. This way, we set $\gamma_{j,t} = \beta_j$. Finally, the parameters for θ_j for the "heuristic trades" z_t^j have been described in the previous section.

In general, the model parameters are unknown, and they need to be estimated from data. Ideally, our data contains the values of the portfolio positions and asset prices, so that we can calculate both B_t and $P_{i,t}$. After that point, however, we find ourselves dealing with a high-dimensional statistical inference problem, a particularly difficult one, considering the functions z_t^j have nonlinear parameters, some of which are continuous, and others discrete. We address the inference problem in Section 3.6.

Since we do not have the data available, we will choose what we believe are "reasonable" parameters, and use them in our simulations and historical backtests. Note that these parameters are chosen mainly for expositional purposes. For more robust conclusions about the effects of our heuristics, these parameters should be based on empirical data.

3.3.1 Default Strategy and Heuristic Weights

We first need to decide on the proper frequency of portfolio decisions in our model. We believe that a monthly frequency is a reasonable compromise between a short time horizon and a long time horizon. This allows us to maintain the computational tractability of investor portfolio performance over a long horizon, while still allowing the investor to make decisions reasonably often. At the same time, if the investor rebalances his portfolio every month, trading costs will have fewer effects than than rebalancing daily or even more frequently.

We will assume a $x_{j,t} = 60\%$ allocation to the risky asset for our default strategy. This corresponds to the classic 60/40 ratio of equities to bonds, still quite a popular strategy among long-term investors. For the heuristic weights, we will consider two situations. The first assumes at most two of the heuristics are active, so that we consider *pairs of heuristics*. This allows us to isolate the effects of the individual heuristics, since only one or two will affect the investor at any given time. Similarly, we can analyze the

Table 3.1: Heuristics Used in Simulations and Backtests

Heuristic	Related behavioral bias	Nature of heuristic	Horizon for past performance (months)	Heuristic lasts for (months)	Threshold
LT Losses	Loss Aversion	Trend-following	12	3	–5% (loss)
LT Gains	Overconfidence	Trend-following	12	3	15% (gain)
ST Losses	Disposition Effect	Contrarian	3	1	–3% (loss)
ST Gains	Disposition Effect	Contrarian	3	1	5% (gain)
Missing Out	Regret	Either	12	3	gain on risky asset threshold: 5%
Recent Decline	Anchoring	Contrarian	24	6	–10% (loss relative to max)
Binary Streak	Gambler's/Hot Hand Fallacy	Contrarian	6	1	# of months of same return sign: 6

This table lists, for each heuristic, the behavioral bias it relates to and the parameters assumed when carrying out simulations and backtests. We also specify whether the heuristic is trend-following or contrarian in nature. LT stands for long-term, ST stands for short-term.

interaction between each pair of heuristics while varying the weighting of the heuristics in a tractable manner. For every pair of heuristics i and j we consider two cases for their weights $(\gamma_{i,t}, \gamma_{j,t})$: $(60\%, 20\%)$ and $(20\%, 60\%)$. If $i = j$, this corresponds to a case when only one heuristic is active, with a weight of $\gamma_{i,t} = 80\%$. The weights for every other heuristic $k \neq i, j$ are set to 0%. With $K = 7$ possible heuristics, we have a total of $K^2 = 49$ different cases.

The second situation occurs when all the heuristics are active. We will use it to analyze how the heuristics interact with each other simultaneously. In that situation, we assign the same weight of $\gamma_{j,t} = 20\%$ to each heuristic.

3.3.2 Heuristic Parameters

In all our simulations we assume fixed parameters for the seven heuristics; this allows for a more consistent comparison across the different cases. In Table 3.1 we specify, for each heuristic, the bias it relates to, the parameters considered, as well as whether the heuristic is contrarian or trend-following in nature. As we will see later, the nature of the heuristic is an important determinant for how the heuristic affects investment performance. We explain our reasoning behind the chosen parameters below.

Long-Term Losses: With our choice of monthly portfolio rebalancing, a one-year horizon is a natural one to measure portfolio losses over a relatively long period. We assume this heuristic lasts for a shorter period, i.e. one quarter. The threshold for losses is -5%, a moderate loss for a portfolio over a one-year period, given that we use a 10% annual return and a 20% annual volatility for our risky asset returns.

Long-Term Gains: This heuristic is similar to the long-term losses heuristic, using the same horizons for past portfolio performance and heuristic duration. The threshold for gains is 15%, a reasonable number for a "good" return on a portfolio given the underlying asset return dynamics described in the previous paragraph.

Short-Term Losses: Here, an investor looks at past performance over a short horizon (3 months), and is reluctant to liquidate his portfolio if he has suffered a loss over the period. We assume this heuristic lasts for a very short period of time (1 month). The loss threshold should also be closer to 0% than the long-term losses heuristic, since we are dealing with a short-term decision. We use a value of -3% for this number.

Short-Term Gains: This heuristic is similar to the short-term losses heuristic. We also assume a three-month horizon for past performance measurement, and one month for the duration of the heuristic. The gain threshold is still close to 0%, but it should be slightly larger than the corresponding value for short-term losses, since the expected return on the asset is positive. We use 5% for the threshold.

Missing Out: The choice of time horizons in this heuristic will be the same as the long-term heuristics considered above. An investor compares the performance of his portfolio over the past year to the performance of the risky asset. If the heuristic is triggered, the investor will be affected by the heuristic for three months. The non-trivial part of specifying this heuristic, however, is setting the proper thresholds of the risky asset returns and the portfolio returns. We tentatively set the threshold for the gain on the risky asset to 10% (close to its expected return), and the threshold for the gain on the portfolio to 2% (an adequate but low return).

Recent Decline: For this heuristic, an investor compares the current price to the highest price "in recent memory." It makes intuitive sense to assume a long horizon for this heuristic, and we use two years. The heuristic is also assumed to last for a rather long period of time, half a year. We set a threshold of -10% on the asset loss in order for an investor to view the asset as "cheap".

Binary Streak: Our final heuristic is the binary streak heuristic. We assume this

Table 3.2: Parameters for Simulated Processes

Return Process	Parameters
Random Walk	mu = 6%; vol = 19%
AR(1) Process, Rho = 20%	mu = 6%; vol = 19%; rho = 20%
AR(1) Process 2, Rho = –20%	mu = 6%; vol = 19%; rho = –20%
MRS Process, S&P 500	(mu1, mu2) = (12%, –17%), (vol1, vol2) = (15%, 34%), P = $\begin{bmatrix} 99\% & 1\% \\ 8\% & 92\% \end{bmatrix}$
MRS Process, Crash Scenario	(mu1, mu2) = (7%, –40%), (vol1, vol2) = (18%, 80%), P = $\begin{bmatrix} 99\% & 1\% \\ 90\% & 10\% \end{bmatrix}$

This table lists the parameters for the return processes used in simulations. *mu* stands for annualized mean, *vol* for volatility, *rho* for serial correlation. P is the transition probability matrix.

heuristic is active over relatively short horizons. We consider an investor looking at the six past monthly returns, and if the vast majority of them (five or more) are of the same sign, the heuristic is triggered. Once triggered, the heuristic lasts for a month.

3.4 Simulation Analysis

We now consider how an investor affected by the heuristics would fare under five different simulated market conditions. We assume a 0% rate of return on the risk-free asset, and consider three generic underlying return processes for the risky asset: a random walk, an AR(1) process, and a Markov regime-switching (MRS) process. For the AR(1) process and the MRS process, we consider two sets of parameter values. We look at five different processes in total, listed in table 3.2.

The first process is a random walk with an annual mean of 6% and volatility of 19%. These parameters are chosen to match the corresponding statistics for the S&P 500 monthly returns over the 1926 - 2014 period. [4] The next two simulated scenarios are AR(1) processes, using the same (unconditional) annual mean and volatility as the random walk. We consider two values for the serial correlation in returns: -20% and 20%. These values of serial correlation may seem a little extreme; however, they were chosen to see more clearly how the effects of the heuristics change for different autocorrelation values. As a reference, we provide the specification for the log returns r_t following an AR(1) process:

[4] S&P 500 returns were obtained from CRSP.

$$r_t \;=\; \mu \;+\; \rho(r_{t-1}-\mu) \;+\sigma\epsilon_t,\ \epsilon_t \sim WN(0,1) \tag{15}$$

where μ is the mean, σ is the volatility, and ρ is the serial correlation coefficient. Setting $\rho = 0\%$ gives a random walk process.

Finally, we analyze two different specifications for the Markov regime-switching process, one simulating an environment that switches between a bull market and a bear market, and another that switches between a "normal" market and stock market crash conditions. We assume that there are two states, and the log-returns in each state are normally distributed. The MRS model is defined as:

$$r_t \;=\; \mu_{i_t} \;+\; \sigma_{i_t}\epsilon_t \;\;,\;\; \epsilon_t \sim WN(0,1) \tag{16}$$

where $i_t \in \{1,2\}$ is the regime indicator which evolves according to a discrete Markov chain with the transition probability matrix P:

$$P \;=\; \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix} \tag{17}$$

so that $p_{ij} = \mathbb{P}(i_{t+1} = j|i_t = i)$. In the bull market regime (without loss of generality, we will assume this is regime 1), the risky asset's returns are distributed as $N(\mu_1, \sigma_1^2)$, and in the bear market regime (regime 2) its returns are distributed as $N(\mu_2, \sigma_2^2)$ where $\mu_2 < \mu_1$.

We estimate the first set of parameters using monthly returns on the S&P 500 over the 1926 - 2014 period. We see that in the "good" regime, there is an annualized mean of 12% and volatility of 15%, while in the "bad" regime, there is a very negative mean of -17%, and high volatility of 34%. The estimated transition probability matrix implies that the good regime occurs 88% of the time, or about 7 out of 8 times.

The second MRS process corresponds to a case where the good regime is a "normal" environment, with returns similar to the S&P 500, while the bad regime represents a stock market crash, one that occurs quite rarely. We specifically assume an annual mean of -40% and a volatility of 80% in the bad regime. We set a transition probability of 1% from the good regime to the bad one, and an unconditional probability of 1% of the bad regime occurring, yielding the transition probability matrix in Table 3.2.

Now that we have chosen the particular specification of the heuristics, their weights, and the underlying returns process, we are ready to simulate investor behavior. We generate 50,000 paths of monthly returns, each path containing 500 observations (corresponding to approximately 40 years). The first 30 observations are used as a “burn-in", during which the investor follows only the default strategy and is not affected by the heuristics. After that burn-in period, the heuristics become relevant. [5] We analyze the performance statistics of the resulting portfolio by calculating them for each return path and then averaging over all the paths.

3.4.1 Simulations with Pairs of Heuristics

We begin by considering the situation where at most two heuristics can be active. We consider every possible pair of heuristics i and j, so that their weights in the investor decision-making model are 60% and 20%, respectively.

Table 3.3 shows the detailed portfolio performance of an investor in one of our five simulated processes. In particular, we look at the situation where the risky asset returns follow an AR(1) process with a positive serial correlation of 20%. Exposure to the heuristics consistently hurts the investor’s Sharpe ratio and hit ratio, with the short-term gains heuristic and the binary streak heuristic being particularly detrimental to performance. This intuitively makes sense: in a strongly trending environment like our example, any gains or positive returns are expected to continue. In such situations, however, getting out of the market (what these two heuristics typically do) is suboptimal. On the other hand, both of these heuristics decrease volatility, since they encourage the investor to be less invested overall. The long-term losses heuristic appears to be the only heuristic that helps to improve maximum drawdown significantly.

Of the other heuristics, the short-term losses heuristic also hurts performance, using the Sharpe ratio and the maximum drawdown as the basis for comparison, because like the earlier examples, it is contrarian in nature. The other three heuristics – long-term gains, missing out, and recent decline – tend to provide higher returns than the buy-and-hold strategy, but also higher volatility.

[5] The burn-in period is required so that there is a history of past asset returns and portfolio balances which are required to generate heuristics going forward.

Pairs of Heuristics Simulations, AR(1) Process, Rho = 20%

Return	LL	LG	SL	SG	M	D	B
LT Losses	2.6%	2.9%	2.5%	2.2%	2.8%	3.2%	2.3%
LT Gains	3.8%	4.2%	4.1%	3.7%	4.2%	4.9%	3.9%
ST Losses	2.8%	3.6%	3.1%	2.5%	3.1%	3.9%	2.7%
ST Gains	2.0%	2.1%	1.8%	2.0%	2.0%	2.3%	1.8%
Missing Out	3.2%	3.8%	3.3%	2.8%	3.4%	4.0%	2.9%
Recent Decline	4.7%	5.1%	4.7%	4.4%	4.7%	4.7%	4.4%
Binary Streak	1.8%	2.3%	2.0%	1.7%	2.1%	2.7%	2.0%
60/40 Strategy	3.4%						

Volatility	LL	LG	SL	SG	M	D	B
LT Losses	9.5%	10.6%	9.9%	8.6%	9.8%	11.0%	8.9%
LT Gains	13.1%	14.1%	14.8%	13.3%	14.1%	16.3%	13.9%
ST Losses	12.2%	14.1%	12.8%	11.7%	12.8%	14.4%	12.0%
ST Gains	8.4%	9.6%	9.9%	9.1%	9.2%	10.9%	9.1%
Missing Out	10.3%	12.4%	11.9%	10.1%	11.1%	13.4%	10.6%
Recent Decline	16.0%	17.0%	16.1%	15.6%	16.0%	16.0%	15.6%
Binary Streak	9.6%	10.8%	10.6%	9.6%	10.3%	11.9%	10.1%
60/40 Strategy	11.1%						

Sharpe Ratio	LL	LG	SL	SG	M	D	B
LT Losses	0.28	0.27	0.26	0.26	0.28	0.29	0.26
LT Gains	0.29	0.30	0.28	0.28	0.30	0.30	0.28
ST Losses	0.23	0.26	0.25	0.21	0.25	0.27	0.23
ST Gains	0.23	0.22	0.18	0.22	0.22	0.21	0.19
Missing Out	0.31	0.30	0.27	0.28	0.31	0.30	0.28
Recent Decline	0.29	0.30	0.29	0.28	0.29	0.29	0.28
Binary Streak	0.19	0.21	0.19	0.18	0.20	0.22	0.20
60/40 Strategy	0.31						

Max Drawdown	LL	LG	SL	SG	M	D	B
LT Losses	-30.4%	-31.9%	-33.7%	-30.0%	-31.0%	-34.5%	-31.5%
LT Gains	-37.0%	-41.3%	-45.8%	-41.0%	-41.3%	-48.2%	-43.2%
ST Losses	-46.8%	-49.1%	-47.9%	-47.8%	-47.9%	-50.4%	-47.6%
ST Gains	-33.6%	-38.8%	-43.3%	-38.0%	-38.4%	-45.0%	-40.3%
Missing Out	-33.6%	-39.2%	-42.8%	-37.4%	-37.7%	-45.2%	-39.7%
Recent Decline	-52.1%	-53.0%	-52.3%	-51.9%	-52.1%	-52.1%	-51.9%
Binary Streak	-41.3%	-43.6%	-45.3%	-42.8%	-43.1%	-46.9%	-42.8%
60/40 Strategy	-37.7%						

Hit Ratio	LL	LG	SL	SG	M	D	B
LT Losses	36.9%	36.4%	39.6%	37.1%	38.6%	52.8%	40.8%
LT Gains	53.3%	53.3%	53.3%	53.3%	53.3%	53.3%	53.3%
ST Losses	53.3%	53.3%	53.3%	53.3%	53.3%	53.3%	53.3%
ST Gains	45.6%	46.6%	43.7%	44.3%	44.5%	49.0%	45.0%
Missing Out	53.3%	53.3%	53.3%	53.3%	53.3%	53.3%	53.3%
Recent Decline	53.3%	53.3%	53.3%	53.3%	53.3%	53.3%	53.3%
Binary Streak	43.3%	46.3%	43.3%	43.3%	44.0%	47.2%	43.3%
60/40 Strategy	53.3%						

Color Code	
(dark green)	Better by 20%
(light green)	Better by 10%
(pink)	Worse by 10%
(red)	Worse by 20%

Table 3.3: Portfolio performance for an investor exposed to at most two heuristics when the log-returns on the risky asset follow an AR(1) process with a 20% serial correlation. For each cell, the heuristic listed in the row was a weight of 60%, while the heuristic listed in the column has a weight of 20%. Results are averaged over 50,000 simulated paths of 500 monthly observations each, with a “burn-in" period of 30 months. Performance metrics are compared to the default 60/40 buy-and-hold strategy. Cells with metrics that differ significantly from the default strategy are color-coded appropriately. The shortened titles in the columns correspond to the same titles as in the rows: *LL* = *LT Losses*, *LG* = *LT Gains*, *SL* = *ST Losses*, *SG* = *ST Gains*, *M* = *Missing Out*, *D* = *Recent Decline*, *B* = *Binary Streak*. The abbreviation *LT* means long-term, and *ST* means short-term.

In Table 3.4, we compare the Sharpe ratios of the investor portfolio for all five of our returns processes. The results for the random walk and the MRS process in a market crash scenario are very similar, with the long-term losses heuristic giving a significantly worse performance relative to the benchmark, and the other heuristics having little effect.

For the AR(1) process with a serial correlation of 20%, the Sharpe ratio decreases in all situations, as we have discussed earlier. On the other hand, for the AR(1) process with a negative serial correlation of −20%, we see significant outperformance from some heuristics – in particular, from the short-term losses heuristic, the short-term gains heuristic, and the binary streak heuristic. This is readily explained by the fact that these heuristics cause an investor to trade in the opposite direction of the movement of past returns, while the returns process in this scenario also exhibits mean reversion. Meanwhile, the long-term losses heuristic gives a much lower Sharpe ratio than the others, as it leads to an investor staying out of the risky asset following a decline in its price.

We note an interesting combination of heuristics that leads to a significant deviation from the benchmark strategy in this process. The combination of the long-term gains heuristic and the missing out heuristic leads to a decrease in the Sharpe ratio of over 10%, while neither of these two heuristics has much effect on its own. This is due to the very large increase in volatility of that specific combination's returns (14.2% compared to just 11.2% for the buy-and-hold strategy). This takes place in situations where an investor stays out of the risky asset and "misses out" on its returns, followed by a significant and persistent investment in the risky asset due to the missing out heuristic, and then a further increase in investment because of the long-term gains heuristic. However, by that point, the investor has typically taken on too much risk, while returns will eventually turn negative due to mean reversion in the underlying process, causing the overall performance to suffer.

The last process we consider is the MRS process calibrated to historical S&P 500 returns. Here, the short-term losses heuristic and the recent decline heuristic are bad for performance, because they are triggered under conditions of large negative returns on the risky asset, which explicitly signify a negative regime and low subsequent returns, while encouraging the investor to hold more of the risky asset, despite the negative regime. The binary streak heuristic also causes the Sharpe ratio to deteriorate relative to the benchmark, since it is triggered in conditions which implicitly identify a good

Pairs of Heuristics Simulations, Sharpe Ratios across Different Processes

Random Walk	LL	LG	SL	SG	M	D	B
LT Losses	0.25	0.24	0.25	0.25	0.25	0.27	0.26
LT Gains	0.28	0.30	0.30	0.29	0.29	0.31	0.30
ST Losses	0.31	0.31	0.32	0.32	0.32	0.32	0.32
ST Gains	0.29	0.30	0.30	0.30	0.30	0.30	0.30
Missing Out	0.31	0.31	0.32	0.32	0.32	0.32	0.32
Recent Decline	0.31	0.31	0.31	0.31	0.31	0.31	0.31
Binary Streak	0.30	0.30	0.31	0.30	0.30	0.31	0.30
60/40 Strategy	0.32						

AR(1), Rho = 20%	LL	LG	SL	SG	M	D	B
LT Losses	0.28	0.27	0.26	0.26	0.28	0.29	0.26
LT Gains	0.29	0.30	0.28	0.28	0.30	0.30	0.28
ST Losses	0.23	0.26	0.25	0.21	0.25	0.27	0.23
ST Gains	0.23	0.22	0.18	0.22	0.22	0.21	0.19
Missing Out	0.31	0.30	0.27	0.28	0.31	0.30	0.28
Recent Decline	0.29	0.30	0.29	0.28	0.29	0.29	0.28
Binary Streak	0.19	0.21	0.19	0.18	0.20	0.22	0.20
60/40 Strategy	0.31						

AR(1), Rho = -20%	LL	LG	SL	SG	M	D	B
LT Losses	0.22	0.20	0.27	0.25	0.22	0.26	0.25
LT Gains	0.28	0.30	0.33	0.32	0.30	0.31	0.31
ST Losses	0.42	0.39	0.41	0.46	0.41	0.38	0.43
ST Gains	0.38	0.40	0.44	0.40	0.40	0.42	0.42
Missing Out	0.32	0.32	0.38	0.38	0.34	0.34	0.36
Recent Decline	0.33	0.33	0.34	0.34	0.34	0.34	0.34
Binary Streak	0.38	0.37	0.41	0.42	0.38	0.37	0.38
60/40 Strategy	0.34						

MRS, S&P 500	LL	LG	SL	SG	M	D	B
LT Losses	0.46	0.44	0.44	0.46	0.46	0.44	0.45
LT Gains	0.41	0.39	0.38	0.39	0.39	0.39	0.38
ST Losses	0.35	0.36	0.35	0.35	0.35	0.36	0.34
ST Gains	0.43	0.39	0.37	0.40	0.40	0.36	0.38
Missing Out	0.44	0.40	0.38	0.41	0.41	0.37	0.39
Recent Decline	0.35	0.37	0.35	0.34	0.35	0.35	0.34
Binary Streak	0.34	0.33	0.31	0.33	0.33	0.31	0.32
60/40 Strategy	0.41						

MRS, Crash Scen.	LL	LG	SL	SG	M	D	B
LT Losses	0.24	0.23	0.25	0.25	0.25	0.27	0.25
LT Gains	0.27	0.28	0.29	0.28	0.28	0.29	0.29
ST Losses	0.30	0.30	0.30	0.30	0.30	0.30	0.30
ST Gains	0.28	0.28	0.28	0.28	0.28	0.29	0.28
Missing Out	0.30	0.30	0.31	0.31	0.31	0.31	0.31
Recent Decline	0.30	0.30	0.30	0.30	0.30	0.30	0.30
Binary Streak	0.28	0.29	0.29	0.29	0.29	0.29	0.29
60/40 Strategy	0.31						

Color Code	
	Better by 20%
	Better by 10%
	Worse by 10%
	Worse by 20%

Table 3.4: Sharpe Ratios of the investor's portfolio for each pair of heuristics and each returns process. We consider a random walk, two AR(1) processes, and two MRS processes – the first with parameters calibrated to historical S&P 500 monthly returns, and the second with parameters chosen to generate rare "crash" scenarios. For each cell, the heuristic listed in the row was a weight of 60%, while the heuristic listed in the column has a weight of 20%. The shortened titles in the columns correspond to the same titles as in the rows: *LL = LT Losses, LG = LT Gains, SL = ST Losses, SG = ST Gains, M = Missing Out, D = Recent Decline, B = Binary Streak.*

or a bad regime (e.g. a high number of subsequent returns that are all negative or all positive) while trading in the direction opposite to the one that an investor would prefer, had he knowledge of the true regime. Finally, the long-term losses heuristic improves performance in this scenario, because it would advise the investor to stay out of the risky asset following large losses.

We conclude that the underlying returns process has a significant effect on the relative performance of our heuristics, whether considered singly or in pairs. This effect can often be explained by whether a heuristic is contrarian or trend-following in nature, and whether the returns process of the risky asset exhibits persistent trending movement or mean reversion. When returns are generated by a regime-switching process, it is important for investor performance if the trigger for the heuristic can implicitly identify the regime, if the trader is subsequently encouraged to take on a larger position in the risky asset, and if that position favors the true underlying regime. Heuristics generally hurt performance as measured by the Sharpe ratio in all cases and for all processes, with the exception of contrarian heuristics for the mean-reverting AR(1) process, and the long-term losses heuristic for the MRS process based on S&P 500 returns.

3.4.2 Simulated Heuristic Correlations

In this section, we analyze the interaction between the heuristics in more detail. Unlike before, we now assume that all the heuristics are active at the same time; the weight assigned to each heuristic is $\gamma_{i,t} = 20\%$. Recall that each heuristic results in an additional allocation z_t^i to the risky asset. (If the heuristic is not active, this additional allocation is simply equal to zero.) We will look at the pairwise correlations $\rho(z_t^i, z_t^j)$ to determine how often the heuristics affect investor allocation in the same direction.

Table 3.5 lists the correlations of the heuristics for each case and each returns process. There is a remarkable similarity in heuristic correlations across different processes – something we did not observe when looking at their relative performance. This is conceptually important, because it allows us to identify the structure of the heuristics from trading data even when the data comes from different assets and different market environments.

Most of these correlations can be explained by the character of the heuristics, that is, whether they are trend-following and extrapolative or contrarian in nature. The long-term losses heuristic and the long-term gains heuristic are both trend-following, and we see that

Correlations between Heuristics across Different Processes

Random Walk	LL	LG	SL	SG	M	D	B
LT Losses		49%	-22%	-18%	0%	-47%	-14%
LT Gains	49%		-20%	-22%	-8%	-52%	-13%
ST Losses	-22%	-20%		54%	1%	20%	17%
ST Gains	-18%	-22%	54%		-1%	12%	16%
Missing Out	0%	-8%	1%	-1%		7%	-2%
Recent Decline	-47%	-52%	20%	12%	7%		12%
Binary Streak	-14%	-13%	17%	16%	-2%	12%	

AR(1), Rho = 20%	LL	LG	SL	SG	M	D	B
LT Losses		44%	-23%	-18%	2%	-54%	-19%
LT Gains	44%		-20%	-22%	-9%	-52%	-18%
ST Losses	-23%	-20%		50%	2%	21%	24%
ST Gains	-18%	-22%	50%		-1%	10%	20%
Missing Out	2%	-9%	2%	-1%		9%	-2%
Recent Decline	-54%	-52%	21%	10%	9%		16%
Binary Streak	-19%	-18%	24%	20%	-2%	16%	

AR(1), Rho = -20%	LL	LG	SL	SG	M	D	B
LT Losses		53%	-21%	-19%	-1%	-40%	-10%
LT Gains	53%		-19%	-22%	-8%	-48%	-10%
ST Losses	-21%	-19%		57%	1%	19%	13%
ST Gains	-19%	-22%	57%		0%	13%	12%
Missing Out	-1%	-8%	1%	0%		5%	-2%
Recent Decline	-40%	-48%	19%	13%	5%		10%
Binary Streak	-10%	-10%	13%	12%	-2%	10%	

MRS, S&P 500	LL	LG	SL	SG	M	D	B
LT Losses		36%	-23%	-12%	-2%	-56%	-13%
LT Gains	36%		-19%	-21%	-7%	-37%	-15%
ST Losses	-23%	-19%		45%	3%	21%	18%
ST Gains	-12%	-21%	45%		0%	5%	15%
Missing Out	-2%	-7%	3%	0%		9%	-1%
Recent Decline	-56%	-37%	21%	5%	9%		13%
Binary Streak	-13%	-15%	18%	15%	-1%	13%	

MRS, Crash Scen.	LL	LG	SL	SG	M	D	B
LT Losses		47%	-22%	-18%	1%	-49%	-13%
LT Gains	47%		-19%	-22%	-8%	-50%	-13%
ST Losses	-22%	-19%		52%	1%	20%	17%
ST Gains	-18%	-22%	52%		0%	11%	16%
Missing Out	1%	-8%	1%	0%		7%	-2%
Recent Decline	-49%	-50%	20%	11%	7%		12%
Binary Streak	-13%	-13%	17%	16%	-2%	12%	

Color Code	
	Above 50%
	Above 25%
	Below -25%
	Below -50%

Table 3.5: Correlations between different behavioral heuristics affecting a portfolio. We assume that each heuristic has a weight of 20%. For every pair of heuristics i and j, we measure the correlation $\rho(z_t^i, z_t^j)$ between the allocations to the portfolio resulting from the heuristics. The shortened titles in the columns correspond to the same titles as in the rows.

they exhibit a strong positive correlation of around 50% across all processes. Examining them in more detail, the long-term gains heuristic leads to a more positive or unchanged allocation to the risky asset, while the long-term losses heuristic usually leads to an unchanged or a more negative allocation. Similarly, the short-term losses heuristic and the short-term gains heuristic also have a strong positive correlation. Finally, every pair of short-term and long-term heuristics has a negative correlation, because the heuristics operate at different time horizons, although it is not as strong as the earlier examples.

The remaining three heuristics are more heterogeneously correlated than the first four. The missing out heuristic has a low correlation with all the other heuristics, since it is the only heuristic determined both by the price behavior of the risky asset and by portfolio performance. The recent decline heuristic has a large negative correlation with both long-term heuristics. This has an intuitive explanation. Following a significant fall in the risky asset price, the investor would be more inclined to add to his position under the recent decline heuristic, rather than to lower it under the long-term losses heuristic, or to leave it unchanged under the long-term gains heuristic. The binary streak heuristic has a relatively low correlation with other heuristics; however, the correlation is negative relative to the long-term heuristics, and positive relative to the short-term heuristics, since the binary streak heuristic is contrarian in nature, while the long-term heuristics are extrapolative, and the short-term heuristics are contrarian.

While the correlation patterns across the different returns processes are in general very similar, there are two notable exceptions. The first exception is the set of correlations in the MRS S&P 500 process between the short-term losses heuristic and the short-term gains heuristic, between the long-term losses heuristic and the long-term gains heuristic, and between the long-term gains heuristic and the recent decline heuristic. These are smaller in magnitude compared to the other four processes. Generally speaking, for all three of these pairs, one heuristic is active only in the good regime in the MRS S&P 500 process, while the other is active only in the bad regime. [6] This dampens the correlations relative to the other returns processes, where both heuristics can be active at the same time.

The other notable exception is the correlation between the long-term losses heuristic and the recent decline heuristic in the AR(1) process with negative serial correlation. The

[6] For example, the long-term losses heuristic, triggered by very poor returns, would usually be triggered in the bad regime for the risky asset.

correlation between this pair is significantly less negative than for the other processes. This is explained by mean reversion in returns in the AR(1) process. The recent decline heuristic is less likely to be triggered at the same time that the long-term losses heuristic is triggered, since this would require large negative returns over both the 24-month and the 12-month periods. However, this is much more likely in the other returns processes, and it is precisely those times when both heuristics are active that lead to a more strongly negative observed correlation between the heuristics in allocation decisions.

To summarize, there is a common correlation structure between allocations influenced by particular heuristics across different returns processes. Trend-following (extrapolative) heuristics tend to be positively correlated to each other, and contrarian heuristics also tend to be positively correlated to each other. We also observe a negative correlation between every pair of trend-following and contrarian heuristics. Particularly important correlations are found between the long-term losses and gains heuristics (positive), between the short-term losses and gains heuristics (also positive), and between the long-term heuristic and the recent decline heuristic (negative).

3.5 Empirical Analysis

We next examine the historical performance of an investor portfolio exposed to these behavioral heuristics. For notational convenience, we will define an *investment strategy* as a pair of assets, a risky asset and a risk-free asset. As before, the default strategy consists of fixed weights of 60% in the risky asset, and 40% in the risk-free one, reset at the start of each month. We consider five different historical investment strategies, spanning different asset classes:

- 60% S&P 500, 40% US Long-Term Corporate Bonds, 1926 - 2014
- 60% S&P 500, 40% US 1-mo. Treasury Bills, 1926 - 2014
- 60% US Long-Term Corporate Bonds, 40% US 1-mo. Treasury Bills, 1926 - 2014
- 60% US Dollar Index, 40% US 1-mo. Treasury Bills, 1990 - 2014
- 60% S&P GSCI Commodity Index, 40% US 1-mo. Treasury Bills, 1970 - 2014

Data for S&P 500 was obtained from the Center for Research in Security Prices (CRSP), and for US Long-Term Corporate Bonds and T-Bills from Ibbotson and Associates. Returns on the US Dollar Index and on GSCI were obtained from Datastream. Note that transactions costs are ignored for all strategies.

As with our earlier simulations, we first consider pairs of heuristics and their corresponding portfolio performance. We then consider the situation when all the heuristics are active, and study the resulting correlations of allocation between the heuristics.

3.5.1 Empirical Results for Pairs of Heuristics

Figure 3.6 shows the historical performance of an investor exposed to pairs of heuristics versus the popular 60/40 equity-bond strategy. In this example, the long-term losses heuristic significantly outperforms the buy-and-hold strategy in the categories of return, volatility, Sharpe ratio, and maximum drawdown. The short-term gains heuristic and the missing out heuristic also improve investor performance, but only if the investor is also exposed to the long-term losses heuristic. In the other cases, investor performance deteriorates, especially if the investor is heavily influenced by the short-term losses heuristic, the recent decline heuristic, or the binary streak heuristic. These three heuristics are contrarian in nature, and given that the underlying risky asset, the S&P 500, exhibits consistent trends over time, with a serial correlation of 8.7%, it is natural that they have a negative affect on investor performance.

In Figure 3.7, we focus on just one performance metric, the Sharpe ratio, and this time consider all five historical investment strategies. Not surprisingly, the performance for the S&P 500 and Bonds strategy is similar to that of the S&P 500 and T-Bills strategy, with the long-term losses heuristic, the short-term gains heuristic, and the missing out heuristic doing well. Most of the heuristics generally do not affect the Bonds and T-Bills strategy. However, the recent decline heuristic strongly hurts performance in this strategy, and the binary streak heuristic also decreases the overall Sharpe ratio, though to a lesser extent. We believe that these differences in performance stem from the low volatility of both bonds and T-Bills, causing the heuristics to be triggered less frequently.

The comparative patterns of Sharpe ratios while trading the GSCI under the influence of the heuristics look surprisingly similar to those of the S&P 500. Once again, the long-term losses heuristic improves investor performance, especially in the presence

Pairs of Heuristics Simulations, S&P 500 and Corporate Bonds, Jan 1926 - Dec 2014

Return	LL	LG	SL	SG	M	D	B
LT Losses	7.7%	8.2%	6.8%	7.1%	7.6%	7.5%	6.5%
LT Gains	6.9%	6.5%	6.6%	6.4%	6.5%	6.3%	5.9%
ST Losses	5.6%	5.9%	5.7%	5.7%	5.7%	5.4%	4.9%
ST Gains	6.1%	5.8%	5.5%	5.8%	5.9%	5.3%	5.4%
Missing Out	6.5%	6.4%	5.9%	5.8%	6.1%	5.7%	5.4%
Recent Decline	5.2%	5.6%	5.3%	5.1%	5.2%	5.2%	4.8%
Binary Streak	4.2%	4.0%	3.8%	3.8%	3.9%	3.6%	3.8%
60/40 Strategy	6.1%						

Volatility	LL	LG	SL	SG	M	D	B
LT Losses	10.6%	11.4%	11.0%	10.3%	10.6%	11.5%	10.6%
LT Gains	13.4%	14.5%	15.6%	14.3%	14.5%	16.5%	14.8%
ST Losses	14.6%	16.0%	15.2%	14.6%	15.2%	16.3%	15.1%
ST Gains	10.6%	12.1%	12.3%	11.2%	11.2%	13.2%	11.6%
Missing Out	11.2%	13.4%	13.7%	11.9%	12.4%	14.8%	12.7%
Recent Decline	17.4%	17.9%	17.5%	17.3%	17.4%	17.4%	17.3%
Binary Streak	12.6%	13.9%	14.3%	13.1%	13.5%	15.0%	13.4%
60/40 Strategy	12.4%						

Sharpe Ratio	LL	LG	SL	SG	M	D	B
LT Losses	0.72	0.72	0.62	0.69	0.72	0.65	0.61
LT Gains	0.52	0.45	0.42	0.45	0.45	0.38	0.40
ST Losses	0.39	0.37	0.37	0.39	0.37	0.33	0.33
ST Gains	0.58	0.48	0.45	0.52	0.53	0.40	0.47
Missing Out	0.58	0.48	0.43	0.49	0.49	0.39	0.42
Recent Decline	0.30	0.31	0.30	0.30	0.30	0.30	0.28
Binary Streak	0.34	0.29	0.26	0.29	0.29	0.24	0.29
60/40 Strategy	0.49						

Max Drawdown	LL	LG	SL	SG	M	D	B
LT Losses	-38.0%	-37.7%	-38.5%	-35.0%	-38.0%	-41.3%	-39.5%
LT Gains	-60.6%	-71.6%	-75.5%	-70.6%	-71.6%	-79.9%	-73.2%
ST Losses	-72.6%	-77.7%	-76.4%	-72.7%	-76.4%	-80.9%	-76.4%
ST Gains	-47.7%	-53.5%	-51.2%	-40.9%	-40.9%	-57.7%	-44.3%
Missing Out	-55.0%	-68.9%	-71.5%	-64.5%	-66.1%	-76.6%	-68.0%
Recent Decline	-84.2%	-85.1%	-84.2%	-84.1%	-84.2%	-84.2%	-84.2%
Binary Streak	-62.4%	-72.5%	-73.3%	-68.7%	-70.5%	-78.1%	-70.0%
60/40 Strategy	-66.1%						

Hit Ratio	LL	LG	SL	SG	M	D	B
LT Losses	62.9%	63.1%	62.6%	62.7%	62.7%	60.6%	62.6%
LT Gains	58.8%	59.0%	58.3%	59.4%	58.9%	58.9%	59.2%
ST Losses	59.2%	58.9%	59.3%	59.7%	59.3%	59.9%	59.7%
ST Gains	60.4%	61.1%	61.0%	61.4%	61.6%	60.3%	62.1%
Missing Out	59.6%	59.7%	59.4%	60.0%	59.7%	59.7%	60.1%
Recent Decline	59.5%	59.3%	59.7%	59.7%	59.5%	59.5%	59.9%
Binary Streak	58.9%	58.5%	59.0%	59.2%	59.2%	58.6%	59.2%
60/40 Strategy	59.7%						

Color Code	
	Better by 20%
	Better by 10%
	Worse by 10%
	Worse by 20%

Table 3.6: Historical performance of an investor employing a 60/40 strategy of equities and bonds, and exposed to at most two heuristics. For each cell, the heuristic listed in the row was a weight of 60%, while the heuristic listed in the column has a weight of 20%. Cells with metrics that differ significantly from the default strategy are color-coded appropriately. The shortened titles in the columns correspond to the same titles as in the rows: *LL = LT Losses, LG = LT Gains, SL = ST Losses, SG = ST Gains, M = Missing Out, D = Recent Decline, B = Binary Streak*. As before, *LT* means Long-Term, and *ST* means short-term.

Sharpe Ratios for Different Asset Classes and Pairs of Heuristics

S&P 500 and Bonds 1926 - 2014	LL	LG	SL	SG	M	D	B
LT Losses	0.72	0.72	0.62	0.69	0.72	0.65	0.61
LT Gains	0.52	0.45	0.42	0.45	0.45	0.38	0.40
ST Losses	0.39	0.37	0.37	0.39	0.37	0.33	0.33
ST Gains	0.58	0.48	0.45	0.52	0.53	0.40	0.47
Missing Out	0.58	0.48	0.43	0.49	0.49	0.39	0.42
Recent Decline	0.30	0.31	0.30	0.30	0.30	0.30	0.28
Binary Streak	0.34	0.29	0.26	0.29	0.29	0.24	0.29
60/40 Strategy	0.49						

S&P 500 and T-Bills 1926 - 2014	LL	LG	SL	SG	M	D	B
LT Losses	0.67	0.58	0.62	0.67	0.64	0.59	0.58
LT Gains	0.48	0.42	0.39	0.40	0.42	0.36	0.36
ST Losses	0.36	0.35	0.35	0.35	0.35	0.32	0.31
ST Gains	0.58	0.49	0.48	0.56	0.56	0.42	0.46
Missing Out	0.54	0.43	0.40	0.44	0.45	0.37	0.39
Recent Decline	0.30	0.31	0.30	0.30	0.30	0.30	0.28
Binary Streak	0.30	0.27	0.25	0.27	0.28	0.24	0.27
60/40 Strategy	0.45						

Bonds and T-Bills 1926 - 2014	LL	LG	SL	SG	M	D	B
LT Losses	1.05	1.02	1.04	1.07	1.05	0.96	1.05
LT Gains	0.99	0.98	0.95	1.00	0.98	0.88	0.97
ST Losses	0.97	0.94	0.97	0.99	0.97	0.88	0.95
ST Gains	1.11	1.09	1.06	1.10	1.10	0.95	1.06
Missing Out	1.06	1.01	1.00	1.07	1.03	0.91	1.01
Recent Decline	0.80	0.81	0.80	0.81	0.80	0.80	0.77
Binary Streak	0.89	0.89	0.87	0.91	0.88	0.78	0.89
60/40 Strategy	1.03						

DXY and T-Bills 1990 - 2014	LL	LG	SL	SG	M	D	B
LT Losses	0.31	0.31	0.29	0.32	0.30	0.17	0.39
LT Gains	0.24	0.24	0.24	0.24	0.24	0.12	0.25
ST Losses	0.22	0.24	0.24	0.24	0.24	0.12	0.24
ST Gains	0.20	0.18	0.19	0.18	0.18	0.14	0.19
Missing Out	0.24	0.24	0.24	0.24	0.24	0.16	0.25
Recent Decline	0.10	0.08	0.09	0.10	0.10	0.10	0.10
Binary Streak	0.29	0.28	0.27	0.26	0.27	0.21	0.27
60/40 Strategy	0.24						

GSCI and T-Bills 1970 - 2014	LL	LG	SL	SG	M	D	B
LT Losses	0.69	0.68	0.66	0.79	0.70	0.64	0.69
LT Gains	0.50	0.46	0.43	0.47	0.46	0.43	0.44
ST Losses	0.39	0.39	0.40	0.38	0.40	0.39	0.38
ST Gains	0.69	0.59	0.56	0.64	0.63	0.46	0.55
Missing Out	0.65	0.52	0.48	0.60	0.57	0.46	0.52
Recent Decline	0.37	0.38	0.37	0.38	0.37	0.37	0.37
Binary Streak	0.44	0.39	0.37	0.44	0.42	0.37	0.41
60/40 Strategy	0.57						

Color Code	
(bright green)	Better by 20%
(light green)	Better by 10%
(pink)	Worse by 10%
(red)	Worse by 20%

Table 3.7: Sharpe Ratios of historical portfolio returns for each pair of heuristics and each investment strategy. We consider 60/40 strategies involving the S&P 500, US Long-Term Corporate Bonds, the US Dollar Index, the S&P GSCI, and US Treasury Bills. For each of the five sub-tables, the default strategy involves holding 60% of the first asset listed in the top left cell of the table, and 40% of the second asset listed in that cell. For each cell in each sub-table, the heuristic listed in the row was a weight of 60%, while the heuristic listed in the column has a weight of 20%.

of the short-term gains heuristic and the missing out heuristic. However, performance deteriorates in other situations, particularly when the recent decline heuristic is involved. This can be explained by trends in the index, since the monthly serial correlation in GSCI monthly returns is 15.8%, even stronger than in the S&P 500.

Finally, the US Dollar Index strategy is positively influenced by the long-term losses heuristic and the binary streak heuristic. This is quite striking, since the long-term losses heuristic is extrapolative in nature, while the binary streak heuristic is contrarian. On the other hand, both the contrarian short-term gains heuristic and the contrarian recent decline heuristic significantly hurt investor performance. This may be caused by different dynamics of the US Dollar Index operating at different time horizons: trending dynamics over horizons longer than one year, and mean reversion over shorter ones. However, it must be kept in mind that the US Dollar Index loses money over the 1990 - 2014 period, so staying out of this risky asset and switching to T-Bills is of benefit to the investor, unlike the other four strategies, where the risky asset outperforms the risk-free one.

To summarize, the long-term losses heuristic and the short-term gains heuristic improve performance for all the historical investment strategies under consideration. While these heuristics are often viewed as irrational from a classical finance perspective, they actually benefit investors in these cases. Other heuristics are generally detrimental to returns, especially the recent decline heuristic and the short-term losses heuristic. As with our earlier simulations, the historical performance of an investor portfolio under the influence of these heuristics can often be explained by the contrarian or extrapolative nature of the heuristics, and whether the underlying risky asset returns “match" these dynamics.

3.5.2 Empirical Heuristic Correlations

As with our earlier simulations, we now study the interactions between heuristics in our historical portfolio returns. Recall that this analysis assumes that all the heuristics are active, each having a weight of $\gamma_{i,t} = 20\%$. The value $\rho(z_t^i, z_t^j)$ measures the correlation between two heuristics, and signifies how often the heuristics cause the investor portfolio to move in the same direction. We find that the patterns of correlation are generally similar to those found in our simulations; however, for some particular pairs, there are substantial differences.

Table 3.8 lists the correlations between the heuristics. The long-term losses heuristic and long-term gains heuristic have a significant positive correlation, as do the short-term heuristics. However, unlike our simulated processes, this correlation is generally low in our historical returns, and particularly low for the 60% Bonds/40% T-Bills and the 60% Dollar Index/40% T-Bills investment strategies. A possible explanation may be that the volatility is lower for these two strategies, causing many periods when both heuristics are inactive, which would dampen the correlation.

We observe a negative correlation between every short-term heuristic and every long-term heuristic. This correlation is particularly low between the long-term losses heuristic and short-term losses heuristic. Furthermore, this is consistent across all asset classes in our investment strategies. Recall that the negative correlation stems from the fact that the short-term heuristics are contrarian in nature, while the long-term heuristics are extrapolative.

The missing out heuristic has a low correlation with other heuristics. In fact, for the investment strategies not involving the S&P 500, this heuristic is never actually active, so its correlations with other heuristics cannot be calculated. The recent decline heuristic has a large negative correlation with both the long-term gains heuristic and the long-term losses heuristic, and an adequate positive correlation with the short-term gains heuristic, similar to what we observed in our simulations.

Finally, the binary streak heuristic in our historical investment strategies has larger correlations than the observed correlations in our earlier simulations. These correlations are positive with respect to the long-term heuristics, negative with respect to the short-term ones, and positive again with respect to the recent decline heuristic. When we look at the historical data, the binary streak heuristic interacts more with the other heuristics in general. There may be a non-obvious dependence between the conditions that give rise to the binary streak heuristic and the other heuristics, given that the dynamics of historical returns must necessarily be more complicated than in our simulated processes.

We briefly discuss how these correlations compare between different investment strategies. Broadly speaking, these patterns are similar to the correlations in performance we observed when looking at pairs of heuristics. Once again, the S&P 500 and Bonds strategy, the S&P 500 and T-Bills strategy, and the GSCI and T-Bills strategy each show a similar picture. Similarly, the patterns of correlation in the US Bonds and T-Bills strategy

Correlations between Heuristics for Different Asset Classes

S&P 500 and Bonds 1926 - 2014	LL	LG	SL	SG	M	D	B
LT Losses		39%	-39%	-8%	-1%	-64%	-34%
LT Gains	39%		-20%	-20%	-6%	-33%	-25%
ST Losses	-39%	-20%		30%	-4%	35%	41%
ST Gains	-8%	-20%	30%		2%	-4%	19%
Missing Out	-1%	-6%	-4%	2%		0%	-2%
Recent Decline	-64%	-33%	35%	-4%	0%		29%
Binary Streak	-34%	-25%	41%	19%	-2%	29%	

S&P 500 and T-Bills 1926 - 2014	LL	LG	SL	SG	M	D	B
LT Losses		38%	-36%	-5%	-7%	-61%	-30%
LT Gains	38%		-18%	-20%	-1%	-32%	-22%
ST Losses	-36%	-18%		25%	1%	38%	42%
ST Gains	-5%	-20%	25%		0%	-9%	15%
Missing Out	-7%	-1%	1%	0%		7%	-4%
Recent Decline	-61%	-32%	38%	-9%	7%		29%
Binary Streak	-30%	-22%	42%	15%	-4%	29%	

Bonds and T-Bills 1926 - 2014	LL	LG	SL	SG	M	D	B
LT Losses		5%	-33%	2%	N/A	-46%	-9%
LT Gains	5%		-1%	-35%	N/A	-2%	-12%
ST Losses	-33%	-1%		5%	N/A	32%	20%
ST Gains	2%	-35%	5%		N/A	-23%	6%
Missing Out	N/A	N/A	N/A	N/A		N/A	N/A
Recent Decline	-46%	-2%	32%	-23%	N/A		22%
Binary Streak	-9%	-12%	20%	6%	N/A	22%	

DXY and T-Bills 1990 - 2014	LL	LG	SL	SG	M	D	B
LT Losses		6%	-41%	-9%	N/A	-60%	-32%
LT Gains	6%		-4%	2%	N/A	-10%	-1%
ST Losses	-41%	-4%		9%	N/A	31%	37%
ST Gains	-9%	2%	9%		N/A	-13%	11%
Missing Out	N/A	N/A	N/A	N/A		N/A	N/A
Recent Decline	-60%	-10%	31%	-13%	N/A		14%
Binary Streak	-32%	-1%	37%	11%	N/A	14%	

GSCI and T-Bills 1970 - 2014	LL	LG	SL	SG	M	D	B
LT Losses		48%	-34%	-18%	N/A	-39%	-31%
LT Gains	48%		-10%	-36%	N/A	-38%	-25%
ST Losses	-34%	-10%		38%	N/A	32%	35%
ST Gains	-18%	-36%	38%		N/A	24%	28%
Missing Out	N/A	N/A	N/A	N/A		N/A	N/A
Recent Decline	-39%	-38%	32%	24%	N/A		31%
Binary Streak	-31%	-25%	35%	28%	N/A	31%	

Color Code	
(bright green)	Above 50%
(light green)	Above 25%
(light red)	Below -25%
(bright red)	Below -50%

Table 3.8: Correlations between different behavioral heuristics affecting a portfolio. For every pair of heuristics i and j, we measure the correlation $\rho(z_t^i, z_t^j)$ between the allocations to the portfolio resulting from the heuristics. Each heuristic is assigned a 20% weight to influence an investor portfolio. For each of the five subtables, the default strategy involves holding 60% of the first asset listed in the top left cell of the table, and 40% of the second asset listed in that cell.

and the Dollar Index and T-Bills strategy closely resemble each other, but depart quite a bit from the patterns found in the first three strategies. This is due to the fact that the volatility for the last two strategies is low, and the dynamics of US T-Bills returns play a more important role than in the strategies trading the the S&P 500 or the GSCI, where these assets' returns overwhelm the returns of US T-Bills.

We find the patterns of correlations are very similar to those we observed in our earlier simulations, with every pair of contrarian heuristics and every pair of trend-following heuristics typically moving the investor portfolio in the same direction. Similarly, every mixed pair of contrarian and trend-following heuristics had a negative correlation with each other when using historical data. Particularly large values of correlations are observed for the S&P 500 and US T-Bills strategy, the S&P 500 and US Bonds strategy, and the GSCI and US T-Bills strategy.

3.6 Model Inference

In this section, we demonstrate a potential approach to estimate the model parameters from trading data. We show that a Metropolis-Hastings algorithm with Gibbs sampling provides relatively accurate estimation performance for several of the parameters, including the default allocation, the allocation noise, and the weights of the heuristics. We first outline the framework for performing the estimation and the proposed distributions for each variable. Next we consider the different cases for the true distributions of the heuristic parameters and active heuristics, and investigate how close our parameter estimates are to the true values.

3.6.1 Modeling Framework

Recall the deterministic version of the investor allocation problem:

$$w_{1,t} = x_{1,t} + \sum_{j=1}^{K} \gamma_{j,t} z_t^j \tag{18}$$

We assume there are two assets, with the first asset being the risky one, so that $w_{1,t}$ is the allocation to the risky asset. The variable $x_{1,t}$ is the default strategy, $\gamma_{j,t}$ is the heuristic weight, and z_t^j is the allocation due to the heuristic. The default strategy and the heuristic weights are assumed to be constant, so that $x_{1,t} = \alpha$ and $\beta_j = \gamma_{j,t}$. The heuristics take on the functional forms described earlier, with allocation z_t^j being defined by parameters θ_j that stay constant throughout the investment horizon.

The "big picture" problem is: given the observed portfolio balances and asset prices for a large sample of investors, how do we estimate the parameters $\alpha, \beta_j, \theta_j$?

We make one enhancement to our deterministic framework by introducing noise into the investment decision. After this modification, the investor allocation becomes:

$$w_{1,t} = x_{1,t} + \sum_{j=1}^{K} \gamma_{j,t} z_t^j + \epsilon_t \tag{19}$$

where $\epsilon_t \sim N(0, \sigma^2)$ is a noise term for an investor; we assume the standard deviation σ remains constant and is another parameter we need to estimate.

While we will investigate the estimation performance for a single investor, we also need to think about the case when there are many investors. There are two reasons this is important. The first is prosaic enough: estimates for one particular investor may not be very accurate. For example, if we obtain 10 years of historical monthly allocations, then we still only have 120 observations, a relatively small number. Furthermore, individual investors usually trade infrequently, making the estimation problem even harder. The second reason for considering a sample of many investors is investigative. We are interested in modeling how investors trade as a group, and we would like to discover this group's cross-sectional distribution of parameters.

We handle multiple investors by assuming each investor i has his own set of parameters $\alpha_i, \beta_{i,j}, \theta_{i,j}$ for $i = 1, 2, \ldots, N$. We also make assumptions on the cross-sectional distribution of these parameters, described below.

Denote by $TN(\mu, \sigma^2, a, b)$ the *truncated normal distribution*, which is the distribution of a normal random variable sampled from $N(\mu, \sigma^2)$, conditional on the variable falling into the range $[a, b]$. Denote by $U(a, b)$ the uniform discrete distribution, restricted to the range $[a, b]$.

We can combine all of the parameters for a particular investor into one vector v_i. The order of these parameters is the same for each investor. For example, the first two parameters can be the default allocation and the allocation noise, followed by a vector of parameters for the long-term losses heuristic and its weight, which is followed by a vector of parameters for the long-term gains heuristic and its weight, etc. The vector includes the continuous parameters, such as the default allocation, the noise, the weight for a particular heuristic, and the threshold for a particular heuristic (e.g. the threshold on losses for the long-term losses heuristic). It also includes the discrete parameters, such as the past horizon over which the investor considers portfolio and/or asset returns, as well as the horizon for the duration of the heuristic. Denote by I the index set of continuous parameters and by J the set of discrete ones.

For every continuous parameter $v_{i,k}$ (with $k \in I$) we assume the truncated normal distribution $TN(\mu_k, \sigma_k^2, a_k, b_k)$. For every discrete parameter $v_{i,l}$ (with $l \in J$) we assume the uniform distribution $U(a_l, b_l)$.

The full set of parameters for a cross-section of investors is the *individual parameters* v_i, for $i = 1, 2, \ldots, n$ and the *sampling parameters*: the mean and standard deviation μ_k, σ_k for the continuous parameters, as well as the bounds a_k, b_k for both the continuous and the discrete parameters.

3.6.2 MCMC Estimation

To estimate our parameters, we employ a particular class of Markov chain Monte Carlo (MCMC) algorithms called the Metropolis-Hastings within Gibbs sampler. We use a "one-at-a-time" component sampling approach for this algorithm. The idea is that we are trying to estimate a multi-dimensional vector x by repeatedly proposing new samples $\widehat{x_{t+1}}$ which differ from the previous sample $\widehat{x_t}$ in exactly one vector component.

The general algorithm works as follows. Assume that we have data z on historical balances and prices for investors, and that we intend to estimate a vector v of parameters. Our aim is to construct a Markov chain $\{v_t\}_{t=1,\ldots,T}$ of samples, which resembles samples from the true distribution of v for large values of t.

We introduce some more notation. Let v_t be the vector of components we have at the end of step t. We assume it has dimension D; define $v_{i,t}$ as the component i of v_t. Define $\pi(\tilde{v}) = f(z|\tilde{v})$ as the likelihood function on the data for parameters $\tilde{v}$. For component $v_{i,t}$ we use the proposal density $q_i(\cdot|v)$ that has support $[a_i, b_i]$; we will outline the densities we use a little later. The algorithm works as follows.

Step 1: Choose initial values v_0 for the vector. Component $v_{i,0}$ is chosen to be an arbitrary real number in the range $[a_i, b_i]$ in the case of continuous components, and an arbitrary integer in the range $[a_i, b_i]$ in the case of discrete components.

Steps 2, ..., T: Given a sample $v_{1,t}, \ldots, v_{D,t}$ we choose the new sample as follows. Define v_{old} to be the vector with component j equal to $v_{j,t}$ for $j \leq i$ and equal to $v_{j+1,t}$ for $j > i$. For $i = 1, \ldots, D$ we propose a new value $w \sim q(\cdot|v_{old})$. Define v_{new} to be the same vector as v_{old}, but that has component i equal to w. Sample a uniform $[0,1]$ random variable U. We set $v_{i,t+1} = w$ if:

$$U < \frac{\pi(v_{new})}{\pi(v_{old})} \times \frac{q(v_{i,t}|v_{new})}{q(w|v_{old})} \tag{20}$$

Otherwise we set $v_{i,t+1} = v_{i,t}$.

After the full path $\{v_t\}_{t=1,\ldots,T}$ is generated, we discard the first M samples as a "burn-in" and average the rest to get our estimates:

$$\widehat{v_t} = \frac{\sum_{t=M+1}^{T} v_{i,t}}{T - M} \tag{21}$$

3.6.3 Single Investor

In the case of a single investor, we observe the historical allocations w_t and the prices of the risky assets p_t. Our goal is to estimate the *investor parameters* x and σ, and the heuristic parameters θ_j.

We assume the parameters are independent of each other. This way, whenever we propose a new value w for a parameter i whose past value is $v_{i,t}$, the conditional distribution only depends on w and $v_{i,t}$, and not on the other parameters. This significantly simplifies our calculations.

The initial values for the parameters are listed in Table 3.19 in the Appendix. These are "reasonable" initial values, deliberately set to be distinct from the true values that we will consider later.

Table 3.19 shows the proposal density for each parameter. In the case of continuous parameters, we use a conditional normal distribution, centered at the previous value of the parameter and with relatively wide bounds. More specifically, if for parameter i, the most recent MCMC sampled value is $v_{i,t}$, then we propose a new value $w \sim TN(v_{i,t}, \sigma_i^2, a_i, b_i)$, where σ_i, a_i, b_i are prespecified. For discrete parameters we use a uniform distribution in a prespecified range, conditional on the outcome being not equal to the previous value.

More specifically, the proposal distribution is $U(a_i, b_i)$ conditional on the fact that the sampled value w is distinct from the last sampled value $v_{i,t}$.

Our choice of prespecified parameters is made by taking into account the nature of the parameters we are estimating. For example, it is natural to assume the plausible range of default allocations is $[0, 100\%]$, while the range of thresholds for the long-term losses heuristic is $[-30\%, 0\%]$.

Each iteration of the algorithm for each parameter involves the calculation of the accept/reject probability in (20). The factor due to the parameter proposal $q(v_{i,t}|v_{new})/q(w|v_{old})$ is quite easy to compute, due to our initial assumption. However, the computation of likelihoods $\pi(v_{new})$ and $\pi(v_{old})$ is more complicated. We present the formula for this likelihood and discuss how we manage to significantly speed up its computation.

Suppose we have a sample of parameters $\widehat{v}$, and we need to compute likelihood π. Following (19), the investor allocation is $w_t = \overline{w}_t + \epsilon_t$, where $\overline{w}_t$ is the deterministic component of allocation and ϵ_t is white noise with variance σ^2. This is assumed to be constrained to $[0, 1]$, so the distribution is $w_t \sim TN(\overline{w}_t, \sigma^2, 0, 1)$. Therefore, if we have estimates $\widehat{w}_t$ of the deterministic component and an estimate $\widehat{\sigma}$ of the allocation noise, the log likelihood function becomes:

$$\log \pi(\widehat{v}) = \sum_{t=1}^{T} \Big[\frac{1}{\widehat{\sigma}\sqrt{2\pi}} \frac{(w_t - \widehat{w}_t)^2}{\widehat{\sigma}^2} \Big/ (\Phi(\frac{1 - w_t}{\widehat{\sigma}}) - \Phi(\frac{0 - w_t}{\widehat{\sigma}}))\Big] \tag{22}$$

where Φ is the cumulative distribution function for a standard normal random variable.

We now turn to the hardest part, which is the calculation of $\widehat{w}_t$. Recall the deterministic framework:

$$\widehat{w}_t = \widehat{x} + \sum_{j=1}^{K} \widehat{\gamma}_j \widehat{z}_t^j \tag{23}$$

To compute the above sum, we require the balances and trades up to time t, as well as the active heuristics up to time t, to identify the active heuristics at time t and their associated trades $\widehat{\gamma}_j \widehat{z}_t^j$. This is a time-consuming exercise, even when the number of trades is as small as one hundred. Without any special "tricks", we would need to repeat this procedure for each parameter, on every run, as many as 100,000 times, estimating only a modest ten parameters over 10,000 MCMC runs.

However, the above computation can be sped up significantly. Because we are only changing one parameter at a time, not all of the components of (23) will change if only one parameter is updated. For example, if only the default strategy $\widehat{x}$ is changed, the allocation due to the heuristics will remain the same, and so does not need to be recomputed. Similarly, if a parameter for one particular heuristic is updated, the allocation due to the other heuristics will remain the same, and we only need to compute the historical allocations due to that particular heuristic. Therefore, if at every step k of the MCMC estimation we store the variables $\widehat{x}$, $\widehat{\gamma}_j$, $\widehat{z}_t^j$ then we need to recompute only one of these variables whenever a new parameter is updated in order to find $\widehat{w}_t$.

To develop an intuition about the performance of the MCMC method, we consider a simple situation where there is only one investor who is affected by a single, known heuristic. We present results only for the long-term losses heuristic; the results are similar for other heuristics. We vary the number of historical observations available to the procedure, as well as the true parameters of the investor.

A brief discussion of the data used in the estimation process is called for. We consider three different sample paths for the prices of the risky asset. The lengths of the sample paths are 1000, 500, and 100. Each sample is generated by simulating returns that follow a random walk with annual mean of 6% and volatility of 19%; these are the same parameters as estimated for the historical S&P 500 returns in Table 3.2. The frequency of returns is assumed to be monthly. The return on the risk-free asset is assumed to be zero. Once the prices are generated, we simulate the investor allocations according to the model 19 and the specified parameters.

There are two caveats to keep in mind. First, we cannot start generating the heuristic allocations immediately, because there are no historical prices and balances to consider. Therefore, we use a "burn-in" period of 30 months, during which the investor allocates according to the default strategy, and no heuristics come into play. Following this period, we generate the full model. When we make our estimates, we assume this burn-in period, and only past this point do we use historical allocations.

For the MCMC run, we draw 20,000 samples, and discard the first 10,000 of them.

The second caveat to keep in mind is that we generate the prices and allocations only once. This makes our results dependent on the particular path generated. However, we feel these results will still give us a good intuition about the performance of the algorithm, and we believe they will not significantly change from the average performance over several independent samples. (Also, at this early stage, we simply want to develop an initial intuition for the process.)

For each of three sample paths, we consider 13 variations for the true parameter values. These are created by taking a default set of true parameters, and perturbing one of the six parameters involved, while keeping the remainder unchanged. We consider two alternative values for each of the six parameters.

Table 3.9 presents the estimation performance when 1000 historical observations are available. We see that our parameter estimates are extremely accurate, with most estimates being within 1% of the true value for the continuous parameters, and being exactly equal to the true value for most discrete parameters. We also see that a significant change in the true value for the default strategy or the heuristic weight leads to significantly worse estimates; changing the true value for other parameters does not have such an effect.

We repeat the same exercise for samples of 500 and 100 historical observations, and report the results in Tables 3.21 and 3.22 in the Appendix. For a sample of 500 observations, our results are still good, although now most continuous parameters estimates are within 5% rather than 1% of the true value. For a sample of 100 observations, however, our estimates are quite poor. The only continuous parameters for which we get adequate estimates (i.e. within 20% of the true value) are the default strategy and the allocation noise. The estimates for other parameters are off by more than 20% in most cases. Table 3.10 summarizes the estimation performance for all 13 variations of each of the three sampled cases.

We conclude that when the number of historical observations is relatively large, we can estimate the parameters for a single heuristic quite well. However, if this number is small, our estimates tend to be poor. Assuming a monthly frequency in our method, 100 observations corresponds to about 8 years, close to the typical sample we observe in our data. Going forward, we will assume that there will be only 100 historical observations available. With a large cross-section of investors, however, our estimates will improve significantly, as we will demonstrate shortly.

Table 3.9: MCMC Estimates for Long-Term Losses Heuristic, Path of Length 1000

Default Strategy						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.4%	20.0%	19.5%	100.0%	95.1%
Allocation Noise	10.0%	10.1%	10.0%	9.6%	10.0%	8.8%
Past Horizon	12.0	12.0	12.0	13.0	12.0	11.0
Loss Threshold	-5.0%	-5.1%	-5.0%	-5.0%	-5.0%	-7.8%
Out Horizon	3.0	3.0	3.0	2.0	3.0	4.0
Heuristic Weight	20.0%	21.1%	20.0%	15.3%	20.0%	14.2%

Allocation Noise						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.4%	60.0%	60.0%	60.0%	59.4%
Allocation Noise	10.0%	10.1%	2.0%	2.1%	20.0%	18.9%
Past Horizon	12.0	12.0	12.0	12.0	12.0	12.0
Loss Threshold	-5.0%	-5.1%	-5.0%	-5.0%	-5.0%	-5.0%
Out Horizon	3.0	3.0	3.0	3.0	3.0	3.0
Heuristic Weight	20.0%	21.1%	20.0%	20.0%	20.0%	18.7%

Past Horizon						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.4%	60.0%	60.0%	60.0%	60.3%
Allocation Noise	10.0%	10.1%	10.0%	10.1%	10.0%	10.0%
Past Horizon	12.0	12.0	6.0	6.0	18.0	18.0
Loss Threshold	-5.0%	-5.1%	-5.0%	-5.0%	-5.0%	-5.0%
Out Horizon	3.0	3.0	3.0	3.0	3.0	3.0
Heuristic Weight	20.0%	21.1%	20.0%	19.3%	20.0%	20.1%

Loss Threshold						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.4%	60.0%	59.8%	60.0%	60.1%
Allocation Noise	10.0%	10.1%	10.0%	10.2%	10.0%	9.9%
Past Horizon	12.0	12.0	12.0	12.0	12.0	12.0
Loss Threshold	-5.0%	-5.1%	-10.0%	-10.0%	-15.0%	-15.0%
Out Horizon	3.0	3.0	3.0	3.0	3.0	3.0
Heuristic Weight	20.0%	21.1%	20.0%	19.7%	20.0%	19.3%

Out Horizon						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.4%	60.0%	59.7%	60.0%	58.4%
Allocation Noise	10.0%	10.1%	10.0%	10.0%	10.0%	10.8%
Past Horizon	12.0	12.0	12.0	12.0	12.0	12.0
Loss Threshold	-5.0%	-5.1%	-5.0%	-5.0%	-5.0%	-4.8%
Out Horizon	3.0	3.0	1.0	1.0	5.0	4.0
Heuristic Weight	20.0%	21.1%	20.0%	19.8%	20.0%	17.9%

Heuristic Weight						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.4%	60.0%	55.6%	60.0%	59.1%
Allocation Noise	10.0%	10.1%	10.0%	14.6%	10.0%	12.5%
Past Horizon	12.0	12.0	12.0	11.0	12.0	11.0
Loss Threshold	-5.0%	-5.1%	-5.0%	-6.4%	-5.0%	-5.0%
Out Horizon	3.0	3.0	3.0	4.0	3.0	4.0
Heuristic Weight	20.0%	21.1%	40.0%	34.3%	60.0%	52.7%

Color Codes: Parameter varied

Continuous Parameter difference:	
	under 1%
	under 5%
	under 10%
	under 20%
	20% or more

Discrete Parameter difference:	
	less than 1
	1 or more
	2 or more

This table presents the MCMC parameter estimates in the case of only one investor, exposed only to the long-term losses heuristic, and when 1000 historical observations are available. There are 13 different subcases for the true parameters. In each of the six subtables, we vary only one parameter (shown in turquoise) while keeping the rest unchanged. We list both the true and the estimated parameters for each case. The estimated parameters are color-coded according to how close they are to the corresponding true values.

Table 3.10: Summary of Estimation Performance for Long-Term Losses Heuristic Only

	% Parameter Estimates		
Path Length	Good	Decent	Poor
1000	69%	23%	8%
500	71%	23%	6%
100	36%	24%	40%

This table shows the performance of our MCMC procedure in estimating the true parameters. We assume only the long-term losses heuristic is active. We use three different sample sizes of historical observations of prices and portfolio balances: 100, 500, and 1000. For each sample size, we consider 13 different subcases for the values of the true parameters; these are listed in Tables 3.9, 3.21, and 3.22. Across all cases and all parameters, we count the proportion of parameters for which the estimate is within 5% of the true value (*good* estimates), for which the estimate is within 20% of the true value (*adequate* estimates), and for which the estimate is more than 20% away from the true value (*poor* estimates); for discrete parameters, an estimate is *good* if we get the parameter exactly right, it is *adequate* if we are off by 1, and it is *poor* if we are off by more than 1.

3.6.4 Multiple Investors

We now extend the estimation procedure to a cross-section of multiple investors. There are two approaches we can use. The first approach involves performing a single MCMC run, during which we simultaneously estimate all of the individual parameters for all investors jointly with the sampling parameters. In the second approach, we estimate individual parameters separately for each investor, and then combine the estimates across investors to obtain the sampling parameters.

We discuss the joint MCMC approach first, where we estimate all the parameters at once. Within each MCMC run, we first update the sampling parameters, and then update the individual parameters for each investor. We assume all the individual parameters are independent of each other, and that they are independent across investors.

Updating Individual Parameters: Updating individual parameters with multiple investors is similar to the single investor case. Suppose we are updating a particular parameter $v_{i,k,t}$, where i is the parameter index, k is the investor index, and t is the MCMC stage. If this is a continuous parameter, then the proposal density is $w \sim TN(v_{i,k,t}, \overline{\sigma}_i^2, a_i, b_i)$, where a_i, b_i are the sampling parameters for the upper and lower bounds and $\overline{\sigma}_i^2$ is prespecified. For a discrete parameter, the proposal density is $w \sim U(a_i, b_i)$, conditional on the fact that the sampled value w is distinct from the last sampled value $v_{i,k,t}$; a_i, b_i are the sampling parameters for the upper and lower bounds.

One change is that the log likelihood function $\log \pi$ is now made up of two components. The first comes from the density of the parameters based on the observed data, shown in (22). The second part comes from the conditional distribution of the individual parameters based on the sampling parameters:

$$\log f_{cond} = \sum_{k=1}^{N} \sum_{i=1}^{D} \log f(v_{i,k,t}|z_i) \tag{24}$$

where N is the number of investors, D is the number of parameters, and $f(v_{i,k,t}|z_i)$ is the density of the individual parameter $v_{i,k,t}$ based on the sampling parameter z_i. Recall that in the case of continuous parameters, $z_i = (\mu_i, \sigma_i^2, a_i, b_i)$ and $v_{i,k,t} \sim TN(\mu_i, \sigma_i^2, a_i, b_i)$, while for discrete parameters $z_i = (a_i, b_i)$ and $v_{i,k,t} \sim U(a_i, b_i)$.

While the above conditional likelihood expression involves many terms when the number of investors is large, only one term changes when one individual parameter is updated. Therefore, as long as we keep track of the previous log likelihood function, computing the new function is straightforward.

Updating Sampling Parameters: Here, we assume the sampling parameters are independent of each other. In Table 3.20 in the Appendix we present the initial values and the proposal densities for each of the continuous parameters. Note that we only include the mean and variance parameters, not the bounds.

Unfortunately, we have not been able to develop a good method for estimating the bounds a_i, b_i. This stems from the fact that these bounds depend on the individual parameters, and have a high likelihood of becoming more extreme over time. For example, the upper bound b_i on a particular parameter $v_{i,k}$ for each individual investor k must satisfy the condition that $b_i \geq v_{i,k}$ for all k. If at least one of the individual parameters $v_{i,k}$ is close to b_i, then whenever we sample a new value for b_i, it can only decrease slightly (as it cannot drop below $v_{i,k}$) but it can increase instead. Therefore, over more runs, the MCMC sampled value $b_{i,t}$ has a tendency to become larger and larger, leading to more imprecision in the estimates.

An exactly analogous issue occurs when estimating the bounds a_i, b_i for the discrete parameters.

To mitigate this phenomenon, we have decided to assume that we know the true values of the bounds for each parameter, and we investigate how well we can estimate the continuous sampling parameters within those bounds. The true values for the bounds are the same as the bounds of the distributions listed in Table 3.19.

In the "separated" MCMC estimation approach, individual parameters for each investor are estimated separately and for each investor, we create a separate MCMC chain $\{v_{i,k,t}\}_{t=1,\ldots,T}$ of parameter estimates. The chain is created using the same initial values and proposal distributions as in Table 3.19. Based on this chain, the estimates of the individual parameters are produced as in (21).

We next turn to the more difficult question, the estimation of the sampling parameters. Once again, we focus on the continuous sampling parameters. The general problem is, for a particular parameter i, we need to estimate its cross-sectional mean μ_i and variance σ_i^2 across investors.

The true model is $v_{i,k} = \mu_i + \eta_{i,k}$ where $\eta_{i,k} \sim N(0, \sigma_i^2)$. We have noisy estimates $v_{i,k,t}$ of $v_{i,k}$, which, under reasonable assumptions, converge to some value $\widehat{v}_{i,k}$. The relation between these is:

$$v_{i,k,t} = \widehat{v}_{i,k} + \xi_{i,k,t};\ \widehat{v}_{i,k} = v_{i,k} + \psi_{i,k}$$

where $\xi_{i,k,t} \sim N(0, \sigma_{i,k}^2(\xi))$ and $\psi_{i,k} \sim N(0, \sigma_{i,k}^2(\psi))$. Therefore, we have:

$$v_{i,k,t} = \mu_i + \xi_{i,k,t} + \psi_{i,k} + \eta_{i,k} \tag{25}$$

We can "remove" the noise $\xi_{i,k,t}$ by averaging across time t. However, we cannot separate the noise $\psi_{i,k}$ from $\eta_{i,k}$, and therefore, our estimates of volatility will be inflated. One possible way to adjust for this inflation is to perform a separate estimation for $\sigma_{i,k}^2(\psi)$, where the true value of the parameter $v_{i,k}$ is fixed, run a large number of MCMC chains to get different samples $\widehat{v}_{i,k}$, and then take their sample variance. However, this is beyond the scope of our study.

Based on the above analysis, our estimators are:

$$\widehat{\mu}_i = \frac{1}{N}\sum_{k=1}^{N}\Big[\frac{1}{T-M}\sum_{t=M+1}^{T} v_{i,k,t}\Big] \tag{26}$$

$$\widehat{\sigma}_i^2 = \frac{1}{N-1}\sum_{k=1}^{N}(\widehat{v}_{i,k} - \overline{v}_i)^2 \tag{27}$$

where:

$$\widehat{v}_{i,k} = \frac{\sum_{t=M+1}^{T} v_{i,k,t}}{T-M}; \; \overline{v}_i = \frac{1}{N}\sum_{k=1}^{N} \widehat{v}_{i,k}$$

We investigate the performance of the above algorithm in estimating the true parameters for a cross-section of investors. We assume that data on 1000 investors is available, with 100 historical observations for each. (These are realistic numbers; on a real dataset, we would likely have even more investors.) We start with the simplest case, when only one heuristic is active, then move on to cases when two heuristics are active, and finally consider the case when all the heuristics are active. We find that the separated MCMC algorithm outperforms the joint algorithm in terms of estimation accuracy. Overall, we consistently produce good estimates of the means for the default strategy and the allocation noise, and adequate estimates of many of the heuristic weights. However, estimating the means for the remaining heuristic weights, the heuristic thresholds, and the volatilities of the sampling distributions is challenging.

Throughout this section, we will focus only on estimating the continuous sampling parameters. We know from our analysis of the case of only one investor that 100 observations are not enough to get good estimates for the individual parameters. This time, however, we have a large number of different investors, with individual parameters drawn from a sampling distribution. We should potentially be able to estimate this distribution quite well.

We first consider the case where only the long-term losses heuristic is present. This is analogous to the single investor case in the previous section. There are six individual parameters for each investor, and of these, four are continuous: the default strategy, the allocation noise, the stop-loss threshold, and the heuristic weight. For each of these, we have four parameters that describe the sampling distribution (a truncated normal): the mean, the volatility, and the lower and upper bounds. As we have seen, estimating the bounds is difficult. However, since their effect on the distribution is small, we concentrate instead on the mean and the volatility. In total, we estimate eight parameters: two parameters for each of the four sampling distributions.

We look at nine different cases for the true parameter values: a default case, and two cases from each of the four sampling distributions, varying the mean for those distributions, while keeping all other parameters unchanged.

Using the joint MCMC estimation approach, Table 3.11 lists the true parameter values with their corresponding estimates. Only the default strategy mean is estimated consistently well, while estimates for the other parameters are predominantly poor. In some cases, however, we are successful in estimating the heuristic weight. Additionally, when the allocation noise is low, our estimates for the loss threshold and the heuristic weight means are very close to the true values.

Table 3.12 looks at the same estimation problem employing a separated MCMC approach. As can be seen, our results are much better. In almost every case, and across all parameters, the estimates for the mean are within 20% of the true values, and in many cases within 5%. However, our estimates for volatility still tend to be poor, and higher than the true distribution volatility. As discussed earlier, this is due to the “extra" noise that comes from the estimation of the individual parameters that cannot be isolated from the cross-sectional volatility in the sampling distribution.

We repeat the same exercise, this time assuming only the long-term gains heuristic is active. Tables 3.23 and 3.24 list the estimation performance of the joint and the separated approaches, respectively. Once again, the separated approach does much better than the joint approach, with most estimates of the mean being within 20% of the true values, and often within 5%.

Based on the above results, we conclude that the separated approach is better than the joint approach, and consequently we employ this method in the following, more complicated cases.

We turn our attention to three cases in which exactly two heuristics are active; in particular, the two long-term heuristics, the two short-term heuristics, and the combination of the long-term losses heuristic and the short-term losses heuristic. In each of these, we look at 13 subcases, varying the mean for each of the six continuous sampling distributions, while keeping the other parameters unchanged.

Tables 3.15 and 3.16 summarize the estimation performance for the means and the volatilities, respectively. The estimates of the means are within 20% of the true values over 50% of the time, and within 5% of the true values over 25% of the time. However, the estimates of the volatility are still predominantly poor. Our estimates are less accurate than the single heuristic cases, not only because we are now estimating more parameters, but also because the heuristics are potentially interacting with one another.

When we compare the estimation performance across the three pairs of heuristics, we find that the algorithm best estimates the parameters for the pair of short-term heuristics. For this pair, our estimates of the means are within 5% of the true values 44% of the time, and within 20% of the true values 79% of the time. This performance is due to the short time horizons of the shorter-term heuristics. Since these heuristics are active over shorter horizons, they tend to be activated more frequently, leading to a greater number of observations of the investor portfolio. Since the true parameters for these heuristics are smaller in absolute value that the others (i.e. for horizon values and heuristic thresholds), it is also less likely that this procedure will markedly misestimate these parameters at the individual investor level.

The worst performing pair of heuristics in our estimation process is the pair consisting of the short-term losses heuristic and the long-term losses heuristic. These heuristics are often triggered at the same time, following a large loss, and since they pull the investor portfolio in opposite directions, [7] they cancel each other out. As a result, we observe an unchanged allocation, even though both heuristics are active, which is detrimental to our estimation performance.

In our final analysis of pairs of active heuristics, we investigate which parameters are estimated more accurately in closer detail. Table 3.13 lists the true parameter values and their corresponding estimates for each of the 13 subcases when both the long-term losses heuristic and the long-term gains heuristic are active. Our focus here is only on the means of these estimates. In almost all cases, the estimates of the default strategy and the allocation noise are within 10% of the true values, with the estimates of the default strategy being slightly more accurate. The estimates for the weights of the two heuristics are also usually adequate. Unfortunately, estimating the threshold values still tends to be problematic. We find similar levels of performance for the other two pairs of heuristics (the pair of short-term heuristics and the pair of losses heuristics).

Finally, we look at the most complicated case, in which all heuristics are active. In the default subcase there are seven heuristics, each with a weight of 20%. We also consider two subcases, one changing the default strategy mean from 60% to 20%, the other changing the allocation noise mean from 10% to 2%. We look at seven additional subcases, where we change the mean of the weight of each of the seven heuristics, one by one, from 20% to 5%, while keeping all other parameters unchanged.

[7] The long-term losses heuristic decreases equity allocation, while the short-term losses heuristic increases it.

Table 3.14 shows the estimation performance for the parameter means across all subcases. In this case, estimation performance is now significantly less accurate than before. We are still able to estimate the default strategy and the allocation noise means quite well. For other parameters, however, our estimates tend to be poor. However, we do have some success in estimating the weights for three of the heuristics, in particular, the weights of the short-term losses heuristic, the short-term gains heuristic, and the missing out heuristic. Since these three heuristics tend to be triggered more frequently, there are more observations where the heuristic weights affect portfolio allocation. We are also able to get accurate estimates for some of the thresholds, namely, for the short-term losses heuristic, the missing out heuristic, and the recent decline heuristic. Our intuition here is that the starting values are reasonably close to the true parameter estimates; for different starting parameter values, these estimates may be less accurate.

In the subcases where a heuristic weight is changed from 20% to 5%, the estimate for the new low weight is significantly different from the true estimate. We hypothesize that when the algorithm initially assumes all heuristics are active, it may have trouble identifying when only a subset of heuristics is active, because it may find an attractive alternative linear combination where all heuristics are present that produces a similar trading behavior.

We note two subcases that highlight the interactions between the heuristics and estimation performance. When the weight of the long-term losses heuristic is reduced to 5%, we have more accurate estimates of the short-term losses heuristic weight and the recent decline heuristic weight. These two heuristics are triggered at the same time as the long-term losses heuristic, but affect the risky asset allocation in the opposite direction. Similarly, when the weight of the short-term gains heuristic is reduced to 5%, we have more accurate estimates of the long-term gains heuristic weight, since these two heuristics are also triggered in similar situations yet have opposite effects on allocation.

We conclude that the separated MCMC approach is able to provide good estimates for some sampling parameters. The means of the default strategy and the allocation noise sampling distributions are estimated consistently well even when all the heuristics are active. The means of the heuristic weights also tend to be estimated well, but less consistently, since these estimates are often poor in the scenario where all heuristics are active. Estimating the means of the thresholds and the volatilities of all sampling distributions tends to be very problematic. Finally, estimates for the heuristics which act over shorter time horizons – the short-term losses heuristic, the short-term gains heuristic, and the binary streak heuristic – are more accurate since these heuristics are triggered more frequently.

Table 3.11: Joint MCMC Estimates for Long-Term Losses Heuristic, Cross-Section of 1000 Investors

Default Strategy						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	59.2%	20.0%	20.9%	100.0%	93.1%
Strategy Cr. Vol	5.0%	6.4%	5.0%	6.5%	5.0%	6.1%
Allocation Noise	10.0%	7.3%	10.0%	5.9%	10.0%	6.3%
Noise Cr. Vol	2.0%	4.3%	2.0%	4.7%	2.0%	4.4%
Loss Threshold	-5.0%	-2.1%	-5.0%	-0.1%	-5.0%	-3.3%
Threshold Cr. Vol	1.0%	8.0%	1.0%	8.2%	1.0%	11.5%
Heuristic Weight	20.0%	18.7%	20.0%	16.3%	20.0%	15.1%
Weight Cr. Vol	3.0%	6.4%	3.0%	8.9%	3.0%	6.1%

Allocation Noise						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	59.2%	60.0%	59.9%	60.0%	56.7%
Strategy Cr. Vol	5.0%	6.4%	5.0%	5.0%	5.0%	7.8%
Allocation Noise	10.0%	7.3%	2.0%	0.0%	20.0%	15.2%
Noise Cr. Vol	2.0%	4.3%	2.0%	3.5%	2.0%	3.6%
Loss Threshold	-5.0%	-2.1%	-5.0%	-5.0%	-5.0%	-2.3%
Threshold Cr. Vol	1.0%	8.0%	1.0%	4.5%	1.0%	10.5%
Heuristic Weight	20.0%	18.7%	20.0%	19.8%	20.0%	16.1%
Weight Cr. Vol	3.0%	6.4%	3.0%	3.7%	3.0%	9.6%

Loss Threshold						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	59.2%	60.0%	59.7%	60.0%	59.8%
Strategy Cr. Vol	5.0%	6.4%	5.0%	6.3%	5.0%	6.2%
Allocation Noise	10.0%	7.3%	10.0%	7.6%	10.0%	7.8%
Noise Cr. Vol	2.0%	4.3%	2.0%	3.4%	2.0%	3.7%
Loss Threshold	-5.0%	-2.1%	-10.0%	-7.5%	-15.0%	-11.0%
Threshold Cr. Vol	1.0%	8.0%	1.0%	6.3%	1.0%	6.5%
Heuristic Weight	20.0%	18.7%	20.0%	19.3%	20.0%	17.8%
Weight Cr. Vol	3.0%	6.4%	3.0%	6.2%	3.0%	7.0%

Heuristic Weight						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	59.2%	60.0%	58.1%	60.0%	56.6%
Strategy Cr. Vol	5.0%	6.4%	5.0%	7.0%	5.0%	6.9%
Allocation Noise	10.0%	7.3%	10.0%	8.0%	10.0%	7.8%
Noise Cr. Vol	2.0%	4.3%	2.0%	3.7%	2.0%	4.4%
Loss Threshold	-5.0%	-2.1%	-5.0%	-3.4%	-5.0%	-5.9%
Threshold Cr. Vol	1.0%	8.0%	1.0%	5.3%	1.0%	4.4%
Heuristic Weight	20.0%	18.7%	40.0%	36.7%	60.0%	48.8%
Weight Cr. Vol	3.0%	6.4%	3.0%	7.7%	3.0%	10.2%

Color Codes:		
	(turquoise)	Parameter varied

Continuous Parameter Difference:	
(green)	under 1%
(yellow)	1% or more
(light orange)	5% or more
(orange)	10% or more
(red)	20% or more

This table presents the MCMC sampling parameter estimates in the 1000 investor case, using 100 historical observations for each investor. We assume investors are influenced only by the long-term losses heuristic. The estimation process uses the joint MCMC approach, in which one MCMC chain is generated for all investors and sampling parameters. We present estimates of the sampling parameter means in the odd-numbered rows, and the volatility in the even-numbered rows. There are 9 different subcases for the true sampling parameters. In each of the four subtables, we vary only one parameter (shown in turquoise) while keeping the rest unchanged. We list both the true and the estimated parameters for each case. The parameter estimates are color-coded according to their closeness to the corresponding true values.

Table 3.12: Separated MCMC Estimates for Long-Term Losses Heuristic, Cross-Section of 1000 Investors

Default Strategy						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	59.1%	20.0%	20.9%	100.0%	91.9%
Strategy Cr. Vol	5.0%	5.8%	5.0%	5.5%	5.0%	3.4%
Allocation Noise	10.0%	10.5%	10.0%	9.4%	10.0%	9.3%
Noise Cr. Vol	2.0%	2.5%	2.0%	2.1%	2.0%	2.1%
Loss Threshold	-5.0%	-5.8%	-5.0%	-5.3%	-5.0%	-8.9%
Threshold Cr. Vol	1.0%	3.5%	1.0%	2.9%	1.0%	5.2%
Heuristic Weight	20.0%	19.5%	20.0%	18.2%	20.0%	16.3%
Weight Cr. Vol	3.0%	5.0%	3.0%	7.9%	3.0%	4.6%

Allocation Noise						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	59.1%	60.0%	60.0%	60.0%	56.9%
Strategy Cr. Vol	5.0%	5.8%	5.0%	5.2%	5.0%	7.0%
Allocation Noise	10.0%	10.5%	2.0%	3.0%	20.0%	20.3%
Noise Cr. Vol	2.0%	2.5%	2.0%	1.7%	2.0%	3.3%
Loss Threshold	-5.0%	-5.8%	-5.0%	-5.1%	-5.0%	-8.5%
Threshold Cr. Vol	1.0%	3.5%	1.0%	2.1%	1.0%	5.2%
Heuristic Weight	20.0%	19.5%	20.0%	19.7%	20.0%	17.1%
Weight Cr. Vol	3.0%	5.0%	3.0%	3.2%	3.0%	7.8%

Loss Threshold						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	59.1%	60.0%	59.9%	60.0%	59.6%
Strategy Cr. Vol	5.0%	5.8%	5.0%	5.6%	5.0%	5.4%
Allocation Noise	10.0%	10.5%	10.0%	10.5%	10.0%	10.7%
Noise Cr. Vol	2.0%	2.5%	2.0%	2.6%	2.0%	2.5%
Loss Threshold	-5.0%	-5.8%	-10.0%	-9.2%	-15.0%	-12.8%
Threshold Cr. Vol	1.0%	3.5%	1.0%	3.6%	1.0%	4.5%
Heuristic Weight	20.0%	19.5%	20.0%	19.2%	20.0%	17.3%
Weight Cr. Vol	3.0%	5.0%	3.0%	5.1%	3.0%	6.0%

Heuristic Weight						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	59.1%	60.0%	59.1%	60.0%	58.2%
Strategy Cr. Vol	5.0%	5.8%	5.0%	5.5%	5.0%	5.6%
Allocation Noise	10.0%	10.5%	10.0%	10.5%	10.0%	9.4%
Noise Cr. Vol	2.0%	2.5%	2.0%	2.7%	2.0%	3.2%
Loss Threshold	-5.0%	-5.8%	-5.0%	-4.8%	-5.0%	-4.9%
Threshold Cr. Vol	1.0%	3.5%	1.0%	1.5%	1.0%	1.7%
Heuristic Weight	20.0%	19.5%	40.0%	38.2%	60.0%	52.9%
Weight Cr. Vol	3.0%	5.0%	3.0%	5.5%	3.0%	6.7%

Color Codes:	
(cyan)	Parameter varied

Continuous Parameter Difference:	
(green)	under 1%
(yellow)	1% or more
(light orange)	5% or more
(orange)	10% or more
(red)	20% or more

This table presents the MCMC sampling parameter estimates in the same manner as Table 3.11, except the estimation process uses the separated MCMC approach. Once again, investors are influenced only by the long-term losses heuristic.

Table 3.13: MCMC Estimates for Long-Term Losses and Long-Term Gains Heuristics

Default Strategy						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	56.8%	20.0%	20.3%	100.0%	83.7%
Allocation Noise	10.0%	10.7%	10.0%	9.2%	10.0%	7.6%
LL Threshold	-5.0%	-6.8%	-5.0%	-5.4%	-5.0%	-11.5%
LL Heuristic Weight	20.0%	17.5%	20.0%	17.5%	20.0%	8.9%
LG Threshold	15.0%	10.0%	15.0%	14.4%	15.0%	5.2%
LG Heuristic Weight	20.0%	19.5%	20.0%	11.9%	20.0%	24.4%

Allocation Noise						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	56.8%	60.0%	59.3%	60.0%	55.1%
Allocation Noise	10.0%	10.7%	2.0%	4.2%	20.0%	19.7%
LL Threshold	-5.0%	-6.8%	-5.0%	-5.4%	-5.0%	-9.1%
LL Heuristic Weight	20.0%	17.5%	20.0%	19.2%	20.0%	15.4%
LG Threshold	15.0%	10.0%	15.0%	10.1%	15.0%	9.7%
LG Heuristic Weight	20.0%	19.5%	20.0%	18.9%	20.0%	18.9%

Long-term Losses Heuristic Threshold						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	56.8%	60.0%	58.2%	60.0%	59.3%
Allocation Noise	10.0%	10.7%	10.0%	10.8%	10.0%	10.9%
LL Threshold	-5.0%	-6.8%	-10.0%	-9.9%	-15.0%	-12.7%
LL Heuristic Weight	20.0%	17.5%	20.0%	18.1%	20.0%	17.8%
LG Threshold	15.0%	10.0%	15.0%	11.0%	15.0%	11.5%
LG Heuristic Weight	20.0%	19.5%	20.0%	19.2%	20.0%	19.0%

Long-term Losses Heuristic Weight						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	56.8%	60.0%	54.2%	60.0%	52.3%
Allocation Noise	10.0%	10.7%	10.0%	11.1%	10.0%	9.8%
LL Threshold	-5.0%	-6.8%	-5.0%	-5.1%	-5.0%	-4.9%
LL Heuristic Weight	20.0%	17.5%	40.0%	33.4%	60.0%	47.8%
LG Threshold	15.0%	10.0%	15.0%	9.1%	15.0%	7.9%
LG Heuristic Weight	20.0%	19.5%	20.0%	21.3%	20.0%	22.4%

Long-term Gains Heuristic Threshold						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	56.8%	60.0%	57.5%	60.0%	55.3%
Allocation Noise	10.0%	10.7%	10.0%	10.4%	10.0%	10.9%
LL Threshold	-5.0%	-6.8%	-5.0%	-7.0%	-5.0%	-6.9%
LL Heuristic Weight	20.0%	17.5%	20.0%	17.9%	20.0%	15.9%
LG Threshold	15.0%	10.0%	10.0%	8.4%	20.0%	10.0%
LG Heuristic Weight	20.0%	19.5%	20.0%	21.1%	20.0%	15.9%

Long-term Gains Heuristic Weight						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	56.8%	60.0%	60.2%	60.0%	60.8%
Allocation Noise	10.0%	10.7%	10.0%	10.6%	10.0%	10.8%
LL Threshold	-5.0%	-6.8%	-5.0%	-6.5%	-5.0%	-7.5%
LL Heuristic Weight	20.0%	17.5%	20.0%	20.0%	20.0%	20.5%
LG Threshold	15.0%	10.0%	15.0%	11.2%	15.0%	11.6%
LG Heuristic Weight	20.0%	19.5%	40.0%	31.4%	60.0%	36.2%

Color Codes:	
(turquoise)	Parameter varied

Continuous Parameter Difference:	
(light green)	under 1%
(yellow)	under 5%
(light orange)	under 10%
(orange)	under 20%
(red)	20% or more

This table presents the MCMC sampling parameter estimates for the 1000 investor case, using 100 historical observations for each investor. We assume the investors are affected by two heuristics: the long-term losses heuristic and the long-term gains heuristic. The separated MCMC approach is used for estimation. We show only the estimates of the means.

There are 13 different cases for the true sampling parameters. In each of the six subtables we vary only one parameter (shown in turquoise) while keeping the rest unchanged. We list both the true and the estimated parameters for each case. *LL* stands for *long-term losses*, and *LG* stands for *long-term gains*.

Table 3.14: MCMC Estimates When All Heuristics Are Active

Parameter	Actual	Est.	Actual	Est.	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	57.4%	20.0%	30.1%	60.0%	58.7%	60.0%	60.2%	60.0%	54.7%
Allocation Noise	10.0%	10.5%	10.0%	11.0%	2.0%	4.4%	10.0%	8.7%	10.0%	10.1%
LL Threshold	-5.0%	-10.3%	-5.0%	-9.0%	-5.0%	-8.8%	-5.0%	-11.0%	-5.0%	-10.7%
LL Heuristic Weight	20.0%	9.8%	20.0%	13.0%	20.0%	11.8%	5.0%	7.9%	20.0%	9.2%
LG Threshold	15.0%	10.8%	15.0%	14.8%	15.0%	10.6%	15.0%	11.4%	15.0%	10.1%
LG Heuristic Weight	20.0%	15.3%	20.0%	10.9%	20.0%	17.6%	20.0%	16.0%	5.0%	10.1%
SL Threshold	-3.0%	-4.1%	-3.0%	-5.8%	-3.0%	-3.7%	-3.0%	-3.6%	-3.0%	-3.9%
SL Heuristic Weight	20.0%	17.3%	20.0%	13.4%	20.0%	18.2%	20.0%	20.7%	20.0%	16.4%
SG Threshold	5.0%	5.0%	5.0%	1.1%	5.0%	5.0%	5.0%	5.1%	5.0%	5.1%
SG Heuristic Weight	20.0%	15.9%	20.0%	19.1%	20.0%	18.0%	20.0%	17.2%	20.0%	17.9%
M Bal Threshold	2.0%	1.6%	2.0%	1.7%	2.0%	1.8%	2.0%	1.6%	2.0%	1.6%
M Prc Threshold	10.0%	12.0%	10.0%	12.1%	10.0%	11.9%	10.0%	11.9%	10.0%	11.9%
M Heuristic Weight	20.0%	12.8%	20.0%	13.1%	20.0%	13.1%	20.0%	12.5%	20.0%	12.9%
D Threshold	-20.0%	-16.4%	-20.0%	-18.0%	-20.0%	-16.9%	-20.0%	-17.4%	-20.0%	-18.9%
D Heuristic Weight	20.0%	14.8%	20.0%	12.3%	20.0%	15.2%	20.0%	19.9%	20.0%	17.0%
B Heuristic Weight	20.0%	17.1%	20.0%	16.2%	20.0%	18.2%	20.0%	17.2%	20.0%	16.5%

Parameter	Actual	Est.	Actual	Est.	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	54.7%	60.0%	60.1%	60.0%	57.0%	60.0%	53.5%	60.0%	56.9%
Allocation Noise	10.0%	11.0%	10.0%	9.9%	10.0%	10.2%	10.0%	11.3%	10.0%	11.2%
LL Threshold	-5.0%	-9.9%	-5.0%	-10.1%	-5.0%	-10.3%	-5.0%	-9.2%	-5.0%	-10.9%
LL Heuristic Weight	20.0%	11.1%	20.0%	10.3%	20.0%	9.7%	20.0%	12.8%	20.0%	10.4%
LG Threshold	15.0%	11.3%	15.0%	10.4%	15.0%	10.9%	15.0%	10.7%	15.0%	10.1%
LG Heuristic Weight	20.0%	15.0%	20.0%	18.0%	20.0%	15.9%	20.0%	17.4%	20.0%	15.4%
SL Threshold	-3.0%	-4.8%	-3.0%	-4.9%	-3.0%	-4.0%	-3.0%	-3.8%	-3.0%	-3.9%
SL Heuristic Weight	5.0%	10.9%	20.0%	15.3%	20.0%	17.4%	20.0%	14.8%	20.0%	15.9%
SG Threshold	5.0%	5.4%	5.0%	5.3%	5.0%	4.9%	5.0%	5.4%	5.0%	4.7%
SG Heuristic Weight	20.0%	13.2%	5.0%	8.5%	20.0%	15.7%	20.0%	15.4%	20.0%	15.1%
M Bal Threshold	2.0%	1.6%	2.0%	1.6%	2.0%	1.7%	2.0%	1.6%	2.0%	1.6%
M Prc Threshold	10.0%	12.2%	10.0%	12.0%	10.0%	12.1%	10.0%	12.2%	10.0%	12.1%
M Heuristic Weight	20.0%	12.6%	20.0%	13.2%	5.0%	13.1%	20.0%	13.3%	20.0%	12.5%
D Threshold	-20.0%	-16.5%	-20.0%	-17.2%	-20.0%	-16.9%	-20.0%	-15.5%	-20.0%	-16.5%
D Heuristic Weight	20.0%	13.2%	20.0%	15.7%	20.0%	15.0%	5.0%	10.2%	20.0%	14.6%
B Heuristic Weight	20.0%	16.0%	20.0%	16.1%	20.0%	17.7%	20.0%	16.6%	5.0%	10.4%

Color Codes: (turquoise) Parameter varied

	Parameter Difference:
(green)	under 1%
(yellow)	under 5%
(light orange)	under 10%
(orange)	under 20%
(red)	20% or more

This table presents the MCMC sampling parameter estimates for the case when all heuristics are active. The separated MCMC approach is used for estimation. Only the estimates of the means are included. There are 10 different subcases for the true sampling parameters considered. In each subcase we vary only one parameter (shown in turquoise) while keeping the rest unchanged. We list both the true and the estimated parameters for each case. *LL* stands for *long-term losses*, and *LG* stands for *long-term gains*, *SL* stands for *short-term losses*, *SG* stands for *short-term gains*, *M* stands for *missing out*, *D* stands for *recent decline*, and *B* stands for *binary streak*.

Table 3.15: Estimation Performance for Sampling Parameter Means for a Cross-Section of Investors, Different Heuristic Cases

Heuristics Active	Estimation Approach	% Parameter Estimates		
		Good	Decent	Poor
Long-term Losses	Joint	25%	28%	47%
Long-term Losses	Separated	50%	42%	8%
Long-term Gains	Joint	25%	25%	50%
Long-term Gains	Separated	42%	31%	28%
LL and LG Biases	Separated	26%	37%	37%
SL and SG Biases	Separated	44%	36%	21%
LL and SL Biases	Separated	26%	28%	46%
All Biases	Separated	11%	33%	57%

This table shows how well our MCMC procedure estimates the means of sampling distributions when historical trading data for a cross-section of investors is available. The first column lists which heuristics are active. The second specifies if a joint or a separated MCMC approach is used for estimation. For each specification of active heuristics, we consider several subcases for the true values of these parameters; these are listed in tables 3.11–3.14 and in tables 3.23–3.24. Across all cases and all parameters, we count the portion of parameters for which the estimate is within 5% of the true value (*good* estimates), for which the estimate is between 5% and 20% of the true value (*adequate* estimates), and for which the estimate is more than 20% away from the true value (*poor* estimates). Note that we consider only continuous sampling parameters here.
LL stands for *long-term losses*, *LG* stands for *long-term gains*, *SL* stands for *short-term losses*, and *SG* stands for *short-term gains*.

Table 3.16: Estimation Performance for Sampling Parameter Volatilities for a Cross-Section of Investors, Different Heuristic Cases

Heuristics Active	Estimation Approach	% Parameter Estimates		
		Good	Decent	Poor
Long-term Losses	Joint	3%	0%	97%
Long-term Losses	Separated	6%	25%	69%
Long-term Gains	Joint	0%	8%	92%
Long-term Gains	Separated	8%	25%	67%
LL and LG Biases	Separated	1%	8%	91%
SL and SG Biases	Separated	3%	4%	94%
LL and SL Biases	Separated	3%	12%	86%
All Biases	Separated	2%	5%	93%

This table lists the estimation performance in the same form as table 3.15, except for estimating the *volatilities* of the sampling distributions for the continuous parameters.

3.7 Predictive Analytics

Our ultimate goal is to construct a model that can be used to predict individual investor behavior. Given this goal, it makes sense to see how the MCMC approach performs in predicting investor allocations based on historical data, and to compare this performance to the more traditional regression-based approaches, such as Grinblatt and Keloharju (2001) and Feng and Seasholes (2005). We find that our approach is usually inferior to regression at predicting allocations at the individual level. For predicting aggregate risky asset allocation, however, the MCMC algorithm tends to outperform regression.

3.7.1 The Prediction Problem

We outline the setting of the prediction problem. Assume we have data on historical asset allocations for 1000 investors, and 150 observations for each investor. This data corresponds to the allocation between a risk-free asset, with a return of 0%, and a risky asset, which we choose to be the S&P 500 from June 2003 to November 2015. The frequency of returns and allocations is monthly, a total of 150 observations. We use the first 100 observations of this data for our models and the last 50 for out-of-sample testing. Figure 3-1 plots the S&P 500 returns over the full period.

For each investor, we predict their allocation to the risky asset for each period in the test sample. The calibrated model applies the data on investor balances and asset returns up until the end of period t to predict the allocation in period $t+1$. For the MCMC approach, the model is estimated using only the training sample (although we consider a more advanced approach later). For the regression, the model is estimated using all allocations up to period t and balances and returns up to period $t-1$, so that as t increases, more historical observations are used.

In many real-world settings, we are interested in the aggregate demand. We are able to calculate the aggregate allocation to the risky asset by summing the total dollar amount allocated to the risky asset across investors and dividing it by the total of investor balances. This allows us to combine our individual allocation predictions to get an aggregate allocation value, in both the MCMC approach and the individual regression approach. We also consider an “aggregate regression" method, which uses only the lagged aggregate allocation and lagged asset returns to predict future aggregate allocation.

3.7.2 Performance Measures

We consider three different possible measures of prediction accuracy. The first measure is the root mean squared error (RMSE) calculated over all investors and all periods in the test sample, comparing the allocation predicted by the model to the actual individual allocation. The second measure performs the same calculation, but compares the prediction to the actual individual allocation without the noise term in the model. Recall from Equation (19) that the allocation in each period is composed of the allocation determined by the default strategy and the heuristics, and a noise term independent of historical data; therefore, comparing the prediction to the allocation without using the noise term should give a clearer idea of the quality of the prediction relative to the true model. The third measure looks at the RMSE at the aggregate level, comparing the predicted allocation to the actual aggregate allocation over all periods in the test sample.

We use the following features in the regression-based approaches for predicting allocation to the risky asset:

- changes in portfolio balance over the past 3, 6, and 12 periods (three variables)
- changes in risky asset price over the past 3, 6, and 12 periods (three variables)
- risky asset allocation in the previous period [8]

More rigorously, if we denote by P_t the price of the risky asset and B_t the portfolio balance at the end of period t, and we assume w_t is the risky asset allocation set at the start of period t, then the regression model is:

$$\begin{aligned} w_{t+1} = \beta_0 + \beta_1(\frac{P_t}{P_{t-3}} - 1) + \beta_2(\frac{P_t}{P_{t-6}} - 1) + \beta_3(\frac{P_t}{P_{t-12}} - 1) + \\ \beta_4(\frac{B_t}{B_{t-3}} - 1) + \beta_5(\frac{B_t}{B_{t-6}} - 1) + \beta_6(\frac{B_t}{B_{t-12}} - 1) + \beta_7 w_t + \epsilon_t \end{aligned} \tag{28}$$

For our comparison, we consider the following ten cases which include active heuristics. The first two are rather simple: we assume that both short-term heuristics are active in

[8] This corresponds to a "lag 1" observation for an AR(1) model. We have considered the case that uses the allocation two periods ago, corresponding to an AR(2) model. The results change very marginally, so we do not include them here.

one case, and that both long-term heuristics are active in the other. As in the estimation section, we assume that the MCMC algorithm is aware that only those two heuristics are active. We then consider the more complicated case where all heuristics are active, each having a weight of 20%. We create seven additional cases by changing the weight for each of the heuristics one by one from 20% to 5%, while leaving the remaining parameters unchanged. Overall, we look at 10 cases in total.

In all of these cases, we use simulated data, where we simulate investor allocations to the S&P 500 and their resulting portfolio balances using the model in Equation (19). It is important to note that the MCMC approach has a "home-court" advantage here, since it already has the correct model specification, unlike the regression approaches under consideration. For more conclusive evidence of the prediction accuracy of the MCMC approach we would need to look at historical trading data; however, this is beyond the scope of the paper.

3.7.3 Predictive Accuracy of Single and Pairwise Heuristics

Table 3.17 lists the prediction accuracy of the different approaches. When only the two short-term heuristics are active, the MCMC approach performs extremely well, yielding only 1.12% RMSE in predicting aggregate allocation, and beating both the individual and the aggregate regression approaches. The MCMC approach also does better than either regression in predicting individual allocations; the RMSE for predicting the actual allocation is 12.68%, which is quite close to the actual noise of 10.04% in the model. This good performance is consistent with our earlier estimation results, where we found that the MCMC approach provides accurate estimates of the short-term heuristics' sampling parameters.

When only the two long-term heuristics are active, the MCMC estimation performance is adequate (a 2.12% RMSE for predicting aggregate allocation); however, this is a worse performance than the regression-based approaches. This can be explained by the fact that the MCMC method is relatively poor at estimating the weights and thresholds for long-term heuristics, resulting in worse predictions than in the short-term heuristic case.

3.7.4 Predictive Accuracy for All Heuristics

We now shift our attention to the out-of-sample prediction performance for the eight cases when all heuristics are active. In most cases, the MCMC-based predictions are more accurate than the regression-based ones, except for the cases where the weights of the short-term heuristics was lowered from 20% to 5%. This makes intuitive sense, because the algorithm estimates the parameters for short-term heuristics very well when they have significant influence on the investor portfolio. When their weight is reduced, the estimates and resulting predictions naturally become less accurate. On the other hand, when the weight of the heuristics acting over longer time horizons is reduced, the MCMC algorithm does better at estimating the short-term heuristic parameters, which produces more accurate predictions overall. This is evident in the fact that the RMSE for the MCMC-based predictions is significantly lower than the RMSE for the regression-based ones in the cases where the weight of the long-term losses heuristic and the recent decline heuristic is reduced.

For robustness, we include the sampling parameter mean estimates for these eight cases in Table 3.25. These results are consistent with our earlier results on the estimation performance, namely that the algorithm does well at estimating the default strategy, the allocation noise, the weights of the short-term heuristics and the binary streak heuristic, as well as a few of the heuristic thresholds. In most cases, the estimates for the other parameters are inaccurate. However, it is important to mention that, even though we have poor estimates for many parameters, our aggregate demand predictions are still very good, with an RMSE ranging from 3.3% to 5.4% across all cases.

When it comes to predicting individual asset allocations, the regression method consistently beats the MCMC one. The RMSE between the predicted and the actual allocation is between 15.3% and 16.2% for regression-based predictions, compared to errors between 16.6% and 19.5% when using MCMC. We believe this larger error is due to extremely poor parameter estimates for a few investors, resulting in highly inaccurate predictions. In the section on estimation, we saw that estimation on 100 historical observations is problematic, even when only one heuristic is active.

As discussed earlier, the regression-based approach expands the estimation window with each new available historical observation, while the MCMC-based approach only uses the training period. Therefore, regressions have the "home-court" advantage of more available data, especially when predicting allocations in periods later in the test sample. For example, for predicting the allocation in period 145, the regression approach uses 144 observations to calibrate the model, while the MCMC approach employs only 100. To understand the possible improvement in prediction, we allow the MCMC approach to expand the estimation window. We use the following “MCMC expanding window" method: after every 5 new historical observations, we recalibrate the model by appending them to the estimation window. For example, we use the first 100 observations to predict the allocations in periods 101–105, the first 105 observations to predict the allocations in periods 106–110, and so on.

The last row of Table 3.17 presents the prediction accuracy of this new approach. We see that the RMSE for predicting aggregate allocation declines significantly from 5.41% to 4.03%, with the latter number substantially beating the regression-based prediction errors of 5.96% and 5.57%. The RMSE for predicting individual allocation also improves, although still not enough to beat individual regressions. We believe that if this approach was extended further to recalibrate the model following every new observation, prediction would improve even more. However, the drawback of the MCMC algorithm is it takes a long time to run. [9] Recalibrating too frequently may be impractical. In practice, one would need to consider the tradeoff between improved accuracy and run-time.

To better visualize our prediction accuracy, Figure 3-2 compares the out-of-sample aggregate allocation predictions of the MCMC approach, the aggregate regression approach, and the MCMC expanding window approach. In the early parts of our test period, from October 2011 to December 2012, the predictions of all three approaches are quite close to the actual allocation, the MCMC-based ones performing the best. By June 2013, all three approaches tend to give excessively high predictions, the predictions of the simple MCMC approach quite poor relative to the other two. After January 2014, the MCMC expanding window approach starts giving extremely good predictions, since by that point it can use the new historical observations to produce a more accurate model. Over the final months of the sample, between February 2015 and November 2015, the aggregate regression approach seems to perform better, although the MCMC expanding window model, once recalibrated, gives very good predictions over the last few months,

[9] The time to run MCMC estimation for all heuristics on a sample of 1000 investors, with 100 observations each, is approximately 90 minutes in Matlab on 16 cores of 2.0 GHz CPUs with 64GB RAM, assuming that 20,000 MCMC runs are performed.

from July 2015 onwards. This repeated improvement in prediction accuracy following model recalibration highlights the importance of expanding the estimation window when using the MCMC method.

We conclude that our MCMC algorithm does well at predicting aggregate investor allocation to the risky asset, especially if we use an expanding estimation window. It tends to beat the regression-based approaches even when all heuristics are active, except for the cases where the short-term heuristics have little effect on allocations in the true model. When it comes to predicting individual allocations, however, the MCMC approach is quite poor and worse than regression-based methods. Therefore, we conclude it is reasonable to employ MCMC methods for modeling aggregate demand, although regression-based methods are still worthwhile to consider, especially considering their speed.

3.8 Conclusion

We have presented a new framework for modeling investor decision-making, where we translate well-known behavioral heuristics into simple systematic investment strategies. We cover seven different heuristics, all of which are motivated by existing examples in the literature, using evidence obtained both in experimental settings and from real-world data. Each heuristic involves a *trigger*, i.e. a set of conditions that must hold for performance on the investor portfolio and the market in order for the heuristic to become active, and thus affect an investor's decision-making process. Each heuristic has an associated *trade*, which defines how the heuristic adjusts the investor's portfolio allocation when active.

Based on simulations of different returns processes and historical data, we find that the effect of a heuristic on portfolio performance is often determined by whether a heuristic is contrarian or trend-following in nature, and how well the underlying returns process favors those dynamics. The long-term losses heuristic and the long-term gains heuristic can both be viewed as trend-following, since they encourage the investor to decrease his allocation following poor returns on the risky asset, and increase it should the risky asset experience good returns. In contrast, the short-term losses heuristic and the short-term gains heuristic are contrarian in nature, as the investor would do the opposite in those situations, except at shorter time horizons. A similar logic argues that the recent decline heuristic and the binary streak heuristic are also contrarian.

The results of our simulated heuristics are as follows. As measured by the Sharpe ratio with respect to the buy-and-hold strategy, behavioral heuristics are detrimental to performance in almost all cases, and for all five processes under consideration: for a random walk process, for two AR(1) processes with serial correlation of 20% and −20%, and for two Markov regime-switching processes, one corresponding to S&P 500 returns, and the other modeling rare crash scenarios. We find only two exceptions to this detrimental performance: the AR(1) process with a serial correlation of −20%, where contrarian heuristics tend to improve performance; and the MRS S&P 500 process, where the long-term losses heuristic gives an improvement over the buy-and-hold strategy.

Considering the interactions between heuristics, we show that trend-following heuristics tend to produce similar directional shifts to the investor portfolio, and these are opposite to the direction contrarian heuristics shift the investor allocation. The correlations between the heuristics are consistent across different returns processes, and are also largely consistent across historical returns applied to different asset classes.

Finally, we looked at the historical performance of an investor's portfolio under the influence of these heuristics. The loss aversion heuristic does very well, improving the performance for all five investment strategies under consideration, in large part because the risky assets in these strategies are indices which exhibit a positive serial correlation in returns, so there is a benefit to a buy-and-hold strategy with an extrapolative heuristic. The contrarian heuristics, on the other hand, are detrimental to performance, especially the short-term losses heuristic, the recent decline heuristic, and the binary streak heuristic. The effects of the heuristics are similar in the strategies that invest in the S&P 500 and the GSCI, but are different in strategies that involve US Corporate Bonds and the US Dollar Index, partly because the volatility of the those asset classes is lower, which means that the heuristics are triggered less frequently.

We have also developed a framework for estimating the parameters of our model. A Metropolis-Hastings algorithm with Gibbs sampling provides relatively accurate sampling parameter estimates using data from a large cross-section of investors. Even though the estimation of parameters at the individual level is problematic due to the lack of necessary observations per investor, aggregating these individual parameters gives reasonably good estimates of the sampling distribution. We obtain accurate estimates for the sampling means of the default strategy, the heuristic weights, and the allocation noise. Also, the estimates of the short-term heuristic parameters are improved because these heuristics are triggered more frequently, and thus there are more historical observations during which they are active.

Finally, we compare the MCMC estimation approach with simple regression-based methods for predicting future investor allocations. The MCMC method tends to do better than regressions at predicting aggregate demand. At the level of the individual investor, however, the MCMC approach underperforms regressions because it badly misestimates the parameters for some investors.

While this is beyond the scope of the paper, it would be very interesting to apply our framework and estimation techniques to real-world data, such as historical trading records. The results could then be used to produce enhanced models of investor trading behavior, and to gain insight into how historical actors would react in various market scenarios. More practically, financial advisors and investors could use systematic behavioral models to identify cases when they themselves are more prone to making irrational decisions, and to avoid these cases should they arise. This would create a more robust investment process for all parties, and a greater chance for them to achieve their long-term financial goals.

Table 3.17: RMSE for Predicting S&P 500 Allocation, Different Approaches

Case for Which Heuristics Are Active		Aggregate Allocation			Individual Allocation				
					Actual Allocation		Proposed Allocation		Allocation Noise
		MCMC	Ind Reg	Agg Reg	MCMC	Ind Reg	MCMC	Ind Reg	
Two Heuristics Active	SL SG Heuristics	1.12%	3.12%	2.55%	12.68%	12.89%	7.77%	8.10%	10.04%
	LL LG Heuristics	2.12%	2.22%	1.61%	14.07%	12.31%	9.89%	7.17%	10.04%
All Heuristics Active	All at 20%	5.41%	5.96%	5.57%	19.54%	16.21%	16.71%	12.84%	10.07%
	LL Heuristic at 5%	4.78%	5.44%	4.90%	18.37%	15.76%	15.51%	12.32%	10.00%
	LG Heuristic at 5%	3.03%	5.74%	4.97%	16.66%	15.39%	13.16%	11.59%	10.19%
	SL Heuristic at 5%	5.32%	5.70%	5.29%	18.74%	15.91%	15.81%	12.40%	10.04%
	SG Heuristic at 5%	7.45%	5.25%	5.08%	19.31%	15.33%	16.49%	11.64%	10.15%
	M Heuristic at 5%	5.07%	5.78%	5.33%	18.95%	15.93%	16.21%	12.61%	10.05%
	D Heuristic at 5%	3.32%	5.93%	5.23%	18.23%	15.96%	15.19%	12.47%	10.09%
	B Heuristic at 5%	3.95%	4.84%	4.69%	19.03%	15.54%	16.13%	11.73%	10.26%
	Exp. Window	4.03%	5.96%	5.57%	17.48%	16.21%	14.38%	12.84%	10.07%

This table compares the prediction accuracy in modeling allocation to the S&P 500 when the true model is the one described in (19). We use monthly returns on the S&P 500 over the June 2003 – November 2015 period, with the first 100 returns used for estimation and the last 50 used for out-of-sample prediction.

Each row corresponds to a different active heuristic case. We look at predictions of individual allocations using the MCMC approach (labeled as *MCMC*) and the individual regression approach (labeled as *Ind Reg*), where we estimate the regression model (28) on all data up to the end of period t to predict the allocation in period $t+1$. We compute the RMSE of the predicted allocation relative to the actual allocation across all investors and all periods in the test sample. As a reference point, in the rightmost column of the table, we show the RMSE of the proposed allocation (without the noise term) relative to the actual allocation. We also compute the RMSE of the predicted allocation relative to the proposed allocation across all investors and all periods. For each case, the cells corresponding to the approach giving the lowest RMSE are highlighted.

We also use three approaches to predict the aggregate allocation to the risky asset. The first two are derived from aggregating individual allocation predictions using the MCMC and individual regression approaches across all investors. The third approach (labeled as *Agg Reg*) estimates the regression model (28) using the aggregate balance of all investors and the lagged aggregate allocation to predict future aggregate allocation. Again, the RMSE of the predicted aggregate allocation relative to the actual one is used for measuring accuracy. For each case, the cells corresponding to the approach giving the lowest RMSE are highlighted.

Ten cases of active heuristic are considered. In the first two cases, only two heuristics are active. In the third case, all heuristics are active and have a weight of 20%. In the remaining seven cases we change the weight of each heuristic, one by one, from 20% to 5% while leaving the other parameters unchanged; the table specifies which heuristic in each case. *LL* stands for *long-term losses*, *LG* stands for *long-term gains*, *SL* stands for *short-term losses*, *SG* stands for *short-term gains*, *M* stands for *missing out*, *D* stands for *recent decline*, and *B* stands for *binary streak*.

Finally, the last row shows the prediction performance in the case where all heuristics have a weight of 20%, but a more advanced MCMC approach is used for estimation. Instead of using only the 100 observations in the training sample, for every 5 new available historical observations, those observations are appended to the estimation window and the model is recalibrated. This way, the MCMC estimation procedure is performed 10 separate times.

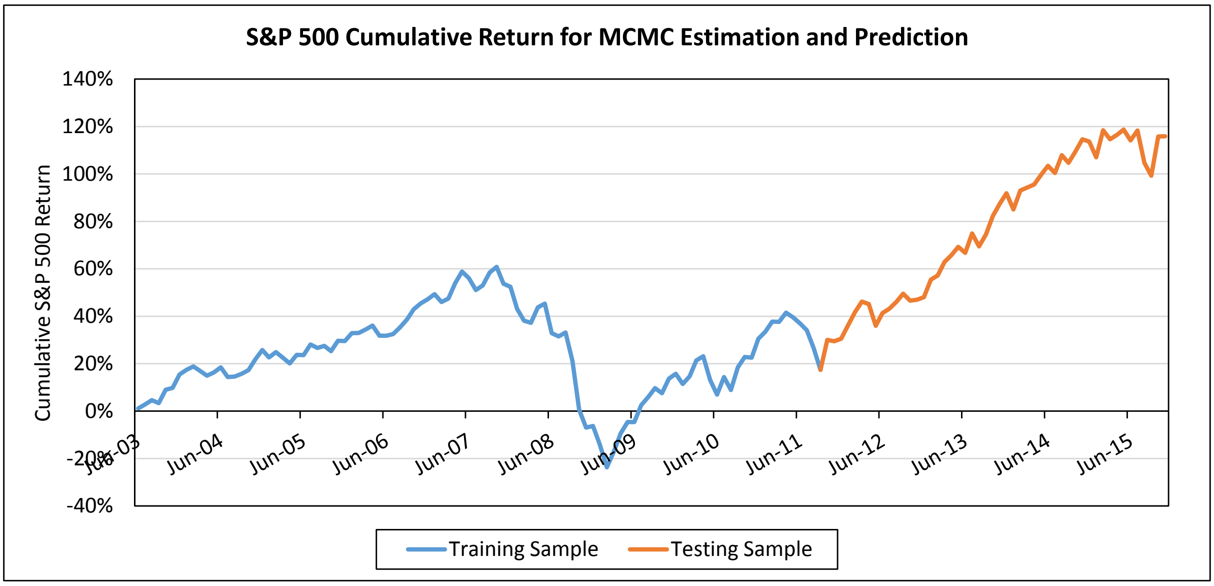

Figure 3-1: Cumulative Returns on the S&P 500 used by the MCMC and regression-based approaches for estimation and prediction.

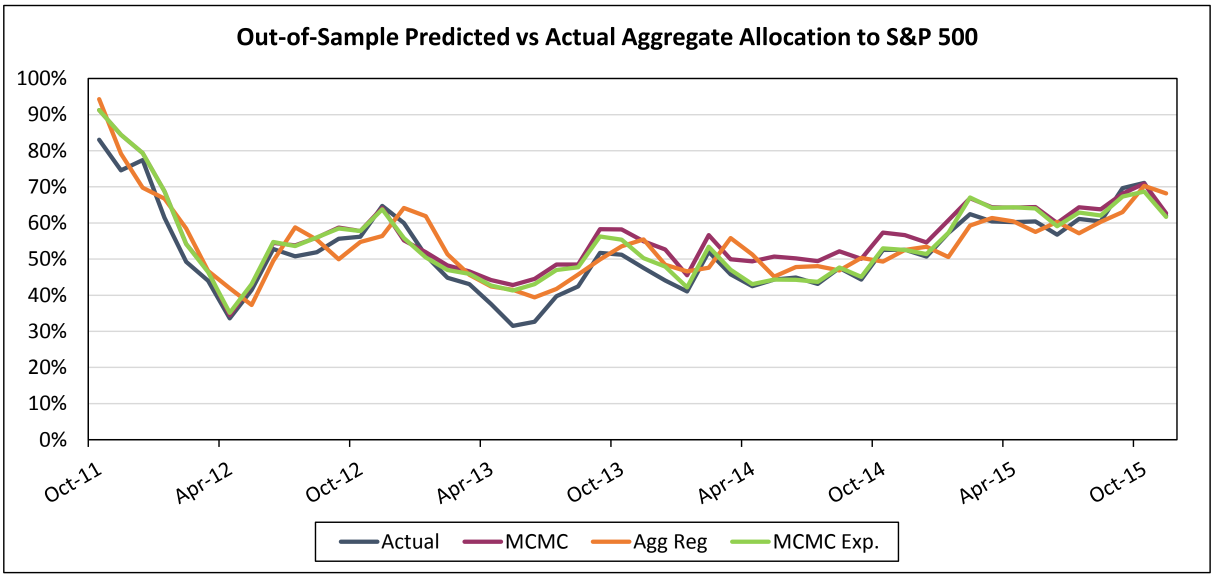

Figure 3-2: This figure plots the actual aggregate allocation to the S&P 500 for the out-of-sample period along with the time series of predicted allocations obtained using three different approaches. *MCMC* is our MCMC approach where the model is estimated on the training set only. *MCMC Exp.* uses the same approach but recalibrates the model by expanding the estimation window every 5 new available historical observations (so that a new estimation is performed on the first 100, 105, 110, etc. observations). *Agg Reg* is the regression approach, where the model (28) is estimated using the aggregate balance of all investors and lagged aggregate allocation. For this model the estimation window is expanded for every new available historical observation.

3.9 Appendix

Sharpe Ratios for One Heuristic, All Processes and Strategies

Heuristic	Simulated Process					Historical Strategy				
	Random Walk	AR(1), Rho = 20%	AR(1), Rho = -20%	MRS, S&P 500	MRS, Crash Scen.	S&P 500 and Bonds 1926 - 2014	S&P 500 and T-Bills 1926 - 2014	Bonds and T-Bills 1926 - 2014	DXY and T-Bills 1990 - 2014	GSCI and T-Bills 1970 - 2014
LT Losses	0.25	0.28	0.22	0.46	0.24	0.72	0.67	1.05	0.31	0.69
LT Gains	0.30	0.30	0.30	0.39	0.28	0.45	0.42	0.98	0.24	0.46
ST Losses	0.32	0.25	0.41	0.35	0.30	0.37	0.35	0.97	0.24	0.40
ST Gains	0.30	0.22	0.40	0.40	0.28	0.52	0.56	1.10	0.18	0.64
Missing Out	0.32	0.31	0.34	0.41	0.31	0.49	0.45	1.03	0.24	0.57
Recent Decline	0.31	0.29	0.34	0.35	0.30	0.30	0.30	0.80	0.10	0.37
Binary Streak	0.30	0.20	0.38	0.32	0.29	0.29	0.27	0.89	0.27	0.41
60/40 Strategy	0.32	0.31	0.34	0.41	0.31	0.49	0.45	1.03	0.24	0.57

Color Code	
	Better by 20%
	Better by 10%
	Worse by 10%
	Worse by 20%

Table 3.18: Sharpe Ratios of portfolio returns when exactly one heuristic is active, with a weight of 80%. We consider five simulated processes and five historical investment strategies. Each cell is color-coded according to how the Sharpe ratio in the cell compares to the Sharpe ratio for the benchmark 60/40 strategy in the same column.

Table 3.19: Initial values and proposal distributions for one investor. For each parameter to be estimated, we specify the parameter, the "group" of parameters to which it belongs, as well as whether it is continuous or not. For the missing out heuristic, *Port. Gain Threshold* is the threshold on the return on the portfolio of the investor, while the *Price Gain Threshold* is the threshold on the return of the risky asset. For the recent decline heuristic, *Price Loss Threshold* is the threshold on the return of the risky asset. For the binary streak heuristic, *Same Sign Thresh.* is the threshold for how many returns of the same signs need to be observed over the past horizon to trigger the heuristic.

Parameter Group	**Parameter**	**Continuous?**	**Initial Value**	**Proposal Density**
Investor Parameters	Default Strategy	Yes	50%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
	Allocation Noise	Yes	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Long-Term Losses	Past Horizon	No	10	$U(\cdot, 6, 18)$
	Loss Threshold	Yes	−7%	$TN(\cdot, 0.1\%, -30\%, 0\%)$
	Active Horizon	No	2	$U(\cdot, 1, 4)$
	Heuristic Weight	Yes	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Long-Term Gains	Past Horizon	No	10	$U(\cdot, 6, 18)$
	Gain Threshold	Yes	7%	$TN(\cdot, 0.1\%, 0\%, 30\%)$
	Active Horizon	No	2	$U(\cdot, 1, 4)$
	Heuristic Weight	Yes	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Short-Term Losses	Past Horizon	No	4	$U(\cdot, 2, 6)$
	Losses Threshold	Yes	−7%	$TN(\cdot, 0.1\%, -15\%, 5\%)$
	Active Horizon	No	2	$U(\cdot, 1, 2)$
	Heuristic Weight	Yes	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Short-Term Gains	Past Horizon	No	4	$U(\cdot, 2, 6)$
	Gains Threshold	Yes	7%	$TN(\cdot, 0.1\%, -5\%, 15\%)$
	Active Horizon	No	2	$U(\cdot, 1, 2)$
	Heuristic Weight	Yes	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Missing Out	Past Horizon	No	10	$U(\cdot, 6, 18)$
	Port. Gain Threshold	Yes	2%	$TN(\cdot, 0.1\%, -2\%, 5\%)$
	Price Gain Threshold	Yes	8%	$TN(\cdot, 0.1\%, 5\%, 20\%)$
	Active Horizon	No	2	$U(\cdot, 1, 4)$
	Heuristic Weight	Yes	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Recent Decline	Past Horizon	No	21	$U(\cdot, 20, 28)$
	Price Loss Threshold	Yes	−15%	$TN(\cdot, 0.1\%, -40\%, 0\%)$
	Active Horizon	No	5	$U(\cdot, 4, 8)$
	Heuristic Weight	Yes	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Binary Streak	Past Horizon	No	7	$U(\cdot, 6, 8)$
	Same Sign Thresh.	No	5	$U(\cdot, 4, 5)$
	Heuristic Weight	Yes	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$

Table 3.20: Initial values and proposal distributions for the continuous sampling parameters. Each parameter corresponds to either the mean or the standard deviation for a particular individual parameter. The parameters corresponding to the same heuristic are grouped together. For the missing out heuristic, *Port. Gain Threshold* is the threshold on the return on the portfolio of the investor, whereas the *Price Gain Threshold* is the threshold on the return of the risky asset. For the recent decline heuristic, *Price Loss Threshold* is the threshold on the return of the risky asset.

Parameter Group	**Parameter**	**Mean/ Std. Dev.?**	**Initial Value**	**Proposal Density**
Investor Parameters	Default Strategy	Mean	50%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
		Std. Dev.	3%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
	Allocation Noise	Mean	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
		Std. Dev.	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Long-Term Losses	Loss Threshold	Mean	−10%	$TN(\cdot, 0.1\%, -30\%, 0\%)$
		Std. Dev.	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
	Heuristic Weight	Mean	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
		Std. Dev.	1%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Long-Term Gains	Gain Threshold	Mean	10%	$TN(\cdot, 0.1\%, 0\%, 30\%)$
		Std. Dev.	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
	Heuristic Weight	Mean	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
		Std. Dev.	1%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Short-Term Losses	Loss Threshold	Mean	−6%	$TN(\cdot, 0.1\%, -15\%, 5\%)$
		Std. Dev.	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
	Heuristic Weight	Mean	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
		Std. Dev.	1%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Short-Term Gains	Gain Threshold	Mean	6%	$TN(\cdot, 0.1\%, 5\%, 15\%)$
		Std. Dev.	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
	Heuristic Weight	Mean	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
		Std. Dev.	1%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Missing Out	Port. Gain Threshold	Mean	3%	$TN(\cdot, 0.1\%, -2\%, 5\%)$
		Std. Dev.	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
	Price Gain Threshold	Mean	9%	$TN(\cdot, 0.1\%, 5\%, 20\%)$
		Std. Dev.	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
	Heuristic Weight	Mean	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
		Std. Dev.	1%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Recent Decline	Price Loss Threshold	Mean	−16%	$TN(\cdot, 0.1\%, -40\%, 0\%)$
		Std. Dev.	5%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
	Heuristic Weight	Mean	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
		Std. Dev.	1%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
Binary Streak	Heuristic Weight	Mean	10%	$TN(\cdot, 0.1\%, 0\%, 100\%)$
		Std. Dev.	1%	$TN(\cdot, 0.1\%, 0\%, 100\%)$

Table 3.21: MCMC Estimates for Long-Term Losses Heuristic, Path of Length 500

Default Strategy

Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.6%	20.0%	20.5%	100.0%	95.7%
Allocation Noise	10.0%	10.5%	10.0%	9.0%	10.0%	8.1%
Past Horizon	12.0	12.0	12.0	14.0	12.0	11.0
Loss Threshold	-5.0%	-5.1%	-5.0%	-4.5%	-5.0%	-7.3%
Out Horizon	3.0	3.0	3.0	1.0	3.0	4.0
Heuristic Weight	20.0%	19.7%	20.0%	15.2%	20.0%	16.4%

Allocation Noise

Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.6%	60.0%	60.1%	60.0%	60.4%
Allocation Noise	10.0%	10.5%	2.0%	2.0%	20.0%	20.2%
Past Horizon	12.0	12.0	12.0	12.0	12.0	11.9
Loss Threshold	-5.0%	-5.1%	-5.0%	-5.1%	-5.0%	-4.5%
Out Horizon	3.0	3.0	3.0	3.0	3.0	2.8
Heuristic Weight	20.0%	19.7%	20.0%	20.2%	20.0%	18.6%

Past Horizon

Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.6%	60.0%	59.5%	60.0%	59.7%
Allocation Noise	10.0%	10.5%	10.0%	10.3%	10.0%	10.3%
Past Horizon	12.0	12.0	6.0	6.0	18.0	18.0
Loss Threshold	-5.0%	-5.1%	-5.0%	-4.9%	-5.0%	-4.8%
Out Horizon	3.0	3.0	3.0	3.0	3.0	3.0
Heuristic Weight	20.0%	19.7%	20.0%	20.0%	20.0%	19.4%

Loss Threshold

Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.6%	60.0%	60.1%	60.0%	60.6%
Allocation Noise	10.0%	10.5%	10.0%	10.5%	10.0%	9.7%
Past Horizon	12.0	12.0	12.0	12.0	12.0	12.0
Loss Threshold	-5.0%	-5.1%	-10.0%	-10.2%	-15.0%	-14.9%
Out Horizon	3.0	3.0	3.0	3.0	3.0	3.0
Heuristic Weight	20.0%	19.7%	20.0%	21.8%	20.0%	23.8%

Out Horizon

Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.6%	60.0%	59.7%	60.0%	58.6%
Allocation Noise	10.0%	10.5%	10.0%	9.7%	10.0%	10.3%
Past Horizon	12.0	12.0	12.0	11.0	12.0	13.2
Loss Threshold	-5.0%	-5.1%	-5.0%	-6.8%	-5.0%	-5.0%
Out Horizon	3.0	3.0	1.0	2.0	5.0	3.4
Heuristic Weight	20.0%	19.7%	20.0%	17.9%	20.0%	18.3%

Heuristic Weight

Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	60.6%	60.0%	60.3%	60.0%	59.3%
Allocation Noise	10.0%	10.5%	10.0%	10.2%	10.0%	9.0%
Past Horizon	12.0	12.0	12.0	12.0	12.0	12.0
Loss Threshold	-5.0%	-5.1%	-5.0%	-4.9%	-5.0%	-5.0%
Out Horizon	3.0	3.0	3.0	3.0	3.0	3.0
Heuristic Weight	20.0%	19.7%	40.0%	41.4%	60.0%	55.1%

Color Codes:

(turquoise)	Param. varied

Continuous parameter difference:	
(green)	under 1%
(yellow)	under 5%
(light orange)	under 10%
(orange)	under 20%
(red)	20% or more

Discrete parameter difference:	
(green)	less than 1
(light orange)	1 or more
(red)	2 or more

This table presents the MCMC parameter estimates for the case when there is only one investor, who is exposed only to the Long-Term Losses heuristic, and 500 historical observations are available. There are 13 different cases for the true parameters. In each of the six subtables we vary only one parameter (shown in turquoise) while keeping the rest unchanged. We list both the true and the estimated parameters for each case. The estimated parameters are color-coded according to how close they are to the corresponding true values.

Table 3.22: MCMC Estimates for Long-Term Losses Heuristic, Path of Length 100

Default Strategy						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	58.6%	20.0%	29.8%	100.0%	92.8%
Allocation Noise	10.0%	10.1%	10.0%	11.5%	10.0%	9.2%
Past Horizon	12.0	15.1	12.0	8.5	12.0	11.8
Loss Threshold	-5.0%	-2.0%	-5.0%	-1.1%	-5.0%	-18.9%
Out Horizon	3.0	2.7	3.0	3.0	3.0	3.5
Heuristic Weight	20.0%	15.8%	20.0%	14.8%	20.0%	12.7%

Allocation Noise						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	58.6%	60.0%	60.1%	60.0%	66.6%
Allocation Noise	10.0%	10.1%	2.0%	2.4%	20.0%	22.2%
Past Horizon	12.0	15.1	12.0	9.8	12.0	11.0
Loss Threshold	-5.0%	-2.0%	-5.0%	-3.7%	-5.0%	-3.7%
Out Horizon	3.0	2.7	3.0	3.0	3.0	2.7
Heuristic Weight	20.0%	15.8%	20.0%	20.0%	20.0%	20.7%

Past Horizon						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	58.6%	60.0%	53.5%	60.0%	59.8%
Allocation Noise	10.0%	10.1%	10.0%	10.5%	10.0%	14.4%
Past Horizon	12.0	15.1	6.0	12.9	18.0	14.7
Loss Threshold	-5.0%	-2.0%	-5.0%	-5.6%	-5.0%	-2.5%
Out Horizon	3.0	2.7	3.0	1.8	3.0	2.9
Heuristic Weight	20.0%	15.8%	20.0%	10.1%	20.0%	19.1%

Loss Threshold						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	58.6%	60.0%	61.5%	60.0%	54.0%
Allocation Noise	10.0%	10.1%	10.0%	10.6%	10.0%	10.0%
Past Horizon	12.0	15.1	12.0	11.9	12.0	12.1
Loss Threshold	-5.0%	-2.0%	-10.0%	-7.6%	-15.0%	-6.9%
Out Horizon	3.0	2.7	3.0	3.0	3.0	1.8
Heuristic Weight	20.0%	15.8%	20.0%	26.9%	20.0%	9.9%

Out Horizon						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	58.6%	60.0%	62.6%	60.0%	54.7%
Allocation Noise	10.0%	10.1%	10.0%	9.7%	10.0%	9.8%
Past Horizon	12.0	15.1	12.0	10.0	12.0	14.5
Loss Threshold	-5.0%	-2.0%	-5.0%	-4.6%	-5.0%	-8.1%
Out Horizon	3.0	2.7	1.0	1.5	5.0	2.9
Heuristic Weight	20.0%	15.8%	20.0%	21.6%	20.0%	12.0%

Heuristic Weight						
Parameter	True	Est.	True	Est.	True	Est.
Default Strategy	60.0%	58.6%	60.0%	57.3%	60.0%	56.9%
Allocation Noise	10.0%	10.1%	10.0%	12.1%	10.0%	7.7%
Past Horizon	12.0	15.1	12.0	11.5	12.0	12.0
Loss Threshold	-5.0%	-2.0%	-5.0%	-3.8%	-5.0%	-5.5%
Out Horizon	3.0	2.7	3.0	2.1	3.0	3.0
Heuristic Weight	20.0%	15.8%	40.0%	30.2%	60.0%	52.9%

Color Codes: Parameter varied

Continuous Parameter difference:	
	under 1%
	under 5%
	under 10%
	under 20%
	20% or more

Discrete Parameter difference:	
	less than 1
	1 or more
	2 or more

This table presents the MCMC parameter estimates for the case when there is only one investor, who is exposed only to the Long-Term Losses heuristic, and 100 historical observations are available. There are 13 different cases for the true parameters. In each of the six subtables we vary only one parameter (shown in turquoise) while keeping the rest unchanged. We list both the true and the estimated parameters for each case. The estimated parameters are color-coded according to how close they are to the corresponding true values.

Table 3.23: Joint MCMC Estimates for Long-Term Gains Heuristic, Cross-Section of 1000 Investors

Default Strategy						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	60.0%	20.0%	19.8%	80.0%	79.1%
Strategy Cr. Vol	5.0%	6.0%	5.0%	5.3%	5.0%	6.0%
Allocation Noise	10.0%	7.5%	10.0%	7.1%	10.0%	5.3%
Noise Cr. Vol	2.0%	4.3%	2.0%	4.3%	2.0%	5.2%
Gain Threshold	15.0%	10.3%	15.0%	12.0%	15.0%	9.2%
Threshold Cr. Vol	1.0%	6.2%	1.0%	11.4%	1.0%	6.9%
Heuristic Weight	20.0%	17.3%	20.0%	1.4%	20.0%	17.9%
Weight Cr. Vol	3.0%	7.3%	3.0%	13.8%	3.0%	9.8%

Allocation Noise						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	60.0%	60.0%	60.2%	60.0%	58.5%
Strategy Cr. Vol	5.0%	6.0%	5.0%	5.3%	5.0%	7.7%
Allocation Noise	10.0%	7.5%	2.0%	0.2%	20.0%	14.6%
Noise Cr. Vol	2.0%	4.3%	2.0%	4.4%	2.0%	3.7%
Gain Threshold	15.0%	10.3%	15.0%	11.0%	15.0%	6.8%
Threshold Cr. Vol	1.0%	6.2%	1.0%	4.6%	1.0%	9.3%
Heuristic Weight	20.0%	17.3%	20.0%	18.0%	20.0%	14.6%
Weight Cr. Vol	3.0%	7.3%	3.0%	4.6%	3.0%	10.3%

Gain Threshold						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	60.0%	60.0%	60.1%	60.0%	59.1%
Strategy Cr. Vol	5.0%	6.0%	5.0%	6.3%	5.0%	6.1%
Allocation Noise	10.0%	7.5%	10.0%	7.3%	10.0%	8.0%
Noise Cr. Vol	2.0%	4.3%	2.0%	4.3%	2.0%	3.4%
Gain Threshold	15.0%	10.3%	10.0%	8.3%	20.0%	9.7%
Threshold Cr. Vol	1.0%	6.2%	1.0%	5.6%	1.0%	9.3%
Heuristic Weight	20.0%	17.3%	20.0%	19.0%	20.0%	8.8%
Weight Cr. Vol	3.0%	7.3%	3.0%	6.5%	3.0%	10.2%

Heuristic Weight						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	60.0%	60.0%	60.6%	60.0%	61.0%
Strategy Cr. Vol	5.0%	6.0%	5.0%	6.1%	5.0%	6.0%
Allocation Noise	10.0%	7.5%	10.0%	8.3%	10.0%	8.4%
Noise Cr. Vol	2.0%	4.3%	2.0%	3.9%	2.0%	3.6%
Gain Threshold	15.0%	10.3%	15.0%	10.9%	15.0%	10.9%
Threshold Cr. Vol	1.0%	6.2%	1.0%	5.6%	1.0%	5.7%
Heuristic Weight	20.0%	17.3%	40.0%	28.8%	60.0%	32.9%
Weight Cr. Vol	3.0%	7.3%	3.0%	9.6%	3.0%	10.3%

Color Codes: Parameter varied

Continuous Parameter Difference:	
(light green)	under 1%
(yellow)	1% or more
(light orange)	5% or more
(orange)	10% or more
(red)	20% or more

This table presents the MCMC sampling parameter estimates for the 1000 investors case, with 100 historical observations for each investor. We assume investors are affected only by the long-term gains heuristic. The joint MCMC approach is used for estimation, where one MCMC chain is generated together for all investors and sampling parameters. We present estimates of the sampling parameter means in the odd-numbered rows, and the volatility in the even-numbered rows.

There are 9 different subcases for the true sampling parameters. In each of the four subtables we vary only one parameter (shown in turquoise) while keeping the rest unchanged. We list both the true and the estimated parameters for each case. The parameter estimates are color-coded according to how close they are to the corresponding true values.

Table 3.24: Separated MCMC Estimates for Long-Term Gains Heuristic, Cross-Section of 1000 Investors

Default Strategy						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	60.0%	20.0%	19.8%	100.0%	93.4%
Strategy Cr. Vol	5.0%	5.6%	5.0%	5.1%	5.0%	2.8%
Allocation Noise	10.0%	10.5%	10.0%	9.9%	10.0%	5.8%
Noise Cr. Vol	2.0%	2.5%	2.0%	2.2%	2.0%	1.8%
Gain Threshold	15.0%	12.2%	15.0%	14.3%	15.0%	10.5%
Threshold Cr. Vol	1.0%	4.6%	1.0%	6.7%	1.0%	5.8%
Heuristic Weight	20.0%	17.9%	20.0%	12.2%	20.0%	22.8%
Weight Cr. Vol	3.0%	6.2%	3.0%	8.4%	3.0%	10.2%

Allocation Noise						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	60.0%	60.0%	60.3%	60.0%	59.4%
Strategy Cr. Vol	5.0%	5.6%	5.0%	5.1%	5.0%	6.8%
Allocation Noise	10.0%	10.5%	2.0%	3.8%	20.0%	19.4%
Noise Cr. Vol	2.0%	2.5%	2.0%	1.7%	2.0%	3.0%
Gain Threshold	15.0%	12.2%	15.0%	11.8%	15.0%	10.8%
Threshold Cr. Vol	1.0%	4.6%	1.0%	3.1%	1.0%	6.0%
Heuristic Weight	20.0%	17.9%	20.0%	19.0%	20.0%	15.9%
Weight Cr. Vol	3.0%	6.2%	3.0%	3.2%	3.0%	7.7%

Gain Threshold						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	60.0%	60.0%	60.4%	60.0%	59.7%
Strategy Cr. Vol	5.0%	5.6%	5.0%	5.6%	5.0%	5.5%
Allocation Noise	10.0%	10.5%	10.0%	10.3%	10.0%	10.8%
Noise Cr. Vol	2.0%	2.5%	2.0%	2.3%	2.0%	2.4%
Gain Threshold	15.0%	12.2%	10.0%	9.3%	20.0%	13.3%
Threshold Cr. Vol	1.0%	4.6%	1.0%	4.0%	1.0%	5.9%
Heuristic Weight	20.0%	17.9%	20.0%	19.0%	20.0%	13.7%
Weight Cr. Vol	3.0%	6.2%	3.0%	5.2%	3.0%	7.3%

Heuristic Weight						
Parameter	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	60.0%	60.0%	60.5%	60.0%	60.8%
Strategy Cr. Vol	5.0%	5.6%	5.0%	5.2%	5.0%	5.5%
Allocation Noise	10.0%	10.5%	10.0%	10.5%	10.0%	10.2%
Noise Cr. Vol	2.0%	2.5%	2.0%	2.7%	2.0%	3.0%
Gain Threshold	15.0%	12.2%	15.0%	12.4%	15.0%	12.4%
Threshold Cr. Vol	1.0%	4.6%	1.0%	4.0%	1.0%	4.0%
Heuristic Weight	20.0%	17.9%	40.0%	32.7%	60.0%	37.2%
Weight Cr. Vol	3.0%	6.2%	3.0%	7.5%	3.0%	8.5%

Color Codes:		Parameter varied

Continuous Parameter Difference:	
	under 1%
	1% or more
	5% or more
	10% or more
	20% or more

This table presents the MCMC sampling parameter estimates in the same way as table 3.23, except when the separated MCMC approach is used for estimation. Again, investors are affected only by the long-term gains heuristic.

Table 3.25: MCMC Estimates When All Heuristics Are Active; S&P 500 Risky Asset

Parameter	Actual	Est.	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	55.2%	60.0%	58.2%	60.0%	53.6%	60.0%	52.7%
Allocation Noise	10.0%	11.3%	10.0%	10.0%	10.0%	11.0%	10.0%	11.4%
LL Threshold	-5.0%	-8.7%	-5.0%	-10.4%	-5.0%	-9.2%	-5.0%	-7.7%
LL Heuristic Weight	20.0%	13.4%	5.0%	8.9%	20.0%	11.7%	20.0%	15.5%
LG Threshold	15.0%	9.2%	15.0%	11.0%	15.0%	9.1%	15.0%	9.3%
LG Heuristic Weight	20.0%	17.8%	20.0%	15.5%	5.0%	12.8%	20.0%	17.4%
SL Threshold	-3.0%	-4.3%	-3.0%	-3.4%	-3.0%	-4.2%	-3.0%	-5.0%
SL Heuristic Weight	20.0%	18.8%	20.0%	21.1%	20.0%	18.2%	5.0%	13.0%
SG Threshold	5.0%	5.7%	5.0%	5.5%	5.0%	5.6%	5.0%	6.1%
SG Heuristic Weight	20.0%	19.5%	20.0%	19.0%	20.0%	19.5%	20.0%	17.5%
M Bal Threshold	2.0%	1.9%	2.0%	1.7%	2.0%	1.8%	2.0%	1.8%
M Prc Threshold	10.0%	11.7%	10.0%	11.9%	10.0%	11.9%	10.0%	11.8%
M Heuristic Weight	20.0%	14.4%	20.0%	13.6%	20.0%	13.3%	20.0%	13.6%
D Threshold	-20.0%	-18.4%	-20.0%	-18.5%	-20.0%	-18.3%	-20.0%	-18.5%
D Heuristic Weight	20.0%	18.5%	20.0%	21.9%	20.0%	18.6%	20.0%	17.0%
B Heuristic Weight	20.0%	15.9%	20.0%	17.2%	20.0%	16.0%	20.0%	13.5%

Parameter	Actual	Est.	Actual	Est.	Actual	Est.	Actual	Est.
Default Strategy	60.0%	57.8%	60.0%	55.5%	60.0%	52.2%	60.0%	55.0%
Allocation Noise	10.0%	10.3%	10.0%	11.1%	10.0%	11.4%	10.0%	11.4%
LL Threshold	-5.0%	-8.4%	-5.0%	-8.2%	-5.0%	-8.7%	-5.0%	-8.6%
LL Heuristic Weight	20.0%	13.3%	20.0%	12.8%	20.0%	15.4%	20.0%	13.3%
LG Threshold	15.0%	9.0%	15.0%	10.1%	15.0%	9.0%	15.0%	9.5%
LG Heuristic Weight	20.0%	18.7%	20.0%	16.8%	20.0%	16.2%	20.0%	18.5%
SL Threshold	-3.0%	-4.7%	-3.0%	-3.9%	-3.0%	-3.1%	-3.0%	-3.8%
SL Heuristic Weight	20.0%	17.0%	20.0%	19.0%	20.0%	16.6%	20.0%	17.4%
SG Threshold	5.0%	7.1%	5.0%	5.6%	5.0%	5.6%	5.0%	6.1%
SG Heuristic Weight	5.0%	13.6%	20.0%	19.1%	20.0%	20.6%	20.0%	17.8%
M Bal Threshold	2.0%	1.7%	2.0%	1.8%	2.0%	2.0%	2.0%	1.7%
M Prc Threshold	10.0%	12.0%	10.0%	11.7%	10.0%	11.6%	10.0%	12.0%
M Heuristic Weight	20.0%	13.1%	5.0%	14.5%	20.0%	14.2%	20.0%	14.0%
D Threshold	-20.0%	-18.1%	-20.0%	-17.9%	-20.0%	-16.1%	-20.0%	-18.7%
D Heuristic Weight	20.0%	19.8%	20.0%	17.9%	5.0%	11.2%	20.0%	19.1%
B Heuristic Weight	20.0%	14.9%	20.0%	16.3%	20.0%	16.4%	5.0%	8.8%

Color Codes:		Param. varied

Parameter Difference:	
(green)	under 1%
(yellow)	under 5%
(light orange)	under 10%
(orange)	under 20%
(red)	20% or more

This table presents the MCMC sampling parameter estimates using historical observations of investor allocations to the S&P 500 over the June 2003 – September 2011 period. All heuristics are active, and eight different cases for sampling parameters are considered. The first is the default case and in the other seven cases we vary only one parameter (shown in turquoise) while keeping the rest unchanged. We list both the true and the estimated parameters for each case. Only estimates for the means are included.

Chapter 4

Preliminary Analysis of Trading Data

(joint work with Daniel Elkind and Andrew W. Lo)

4.1 Introduction

A crucial ingredient for building models of investor behavior is good trading data. Unfortunately, this data was obtained very late during my Ph.D. studies and so we can only include preliminary results here. The results pertain to understanding the overall trading distribution at the security and the portfolio level, as well as how this distribution depends on past performance and trading. We replicate the phenomenon documented by Ben-David and Hirshleifer (2012) whereby the probability that at investor trades increases in the past returns on the investor's positions. We obtain further insights into this result by looking at what trades were made in the past as well varying the lengths of the horizons over which we measure returns and trading.

The dataset we are using is very good. It includes trading records in over 1,000,000 accounts at a large U.S. retail broker over the 2003–2015 period. We have records of all trades made by these accounts, as well as snapshots of positions and flows into the accounts at monthly frequency. The large number of investors in the dataset would enable us to apply machine-learning methods that tend to require large amounts of data. Furthermore, the sample period of 13 years is a long one; we have not found other studies in the literature that have granular trading data available over such a long timespan. The advantage of a long sample period is that it incorporates different market environments so that we are able to analyze how people behave during various trading conditions.

We briefly review the major papers that have also analyzed investor trading data. Barber and Odean have carried out comprehensive studies of individual investors demonstrating the disposition effect (Odean, 1998), poor trading performance (Barber and Odean, 2000), more frequent trading by men (Barber and Odean, 2001), increased likelihood of buying attention-grabbing stocks (Barber and Odean, 2008), and reluctance to buy stocks for which investors previously lost money and/or missed out on gains (Strahilevitz, Odean, and Barber, 2011). Grinblatt and Keloharju have studied the detailed Finnish dataset; they identify different investor types (Grinblatt and Keloharju, 2000), investigate the effects of familiarity in terms of distance, language, and culture of firms on the likelihood of trading their shares by investors (Grinblatt and Keloharju, 2001a), and find increased propensity to trade by overconfident and "sensation-seeking" investors (Grinblatt and Keloharju, 2009). While these papers document important features of how people trade, they do not address the question of how these features interact with each other and to which extent they allow for predicting future trading decisions.

A few studies have employed a regression approach to address the question of predicting trading. Grinblatt and Keloharju (2001b) were the first to conduct a major study of how investor trading is related to past returns by running logit regressions of the hold/sell indicator variable on values of past returns, as well as on investor characteristics. They also analyze the sell vs. buy decision in similar manner. It is important to note that they restrict their observation days to only those when an investor buys or sells a stock. While this approach gives valuable insights on how trading in an investor's portfolio on those days relates to past features, it does not account for situations where these features (such as large losses or large gains) were present, but no subsequent trading took place. Therefore, the method is not enough to predict future trading behavior.

Feng and Seasholes (2005) address this drawback by modeling the survival of stock holdings in the investor's portfolio. They estimate the "survival" function for how long a stock is held before being sold, as well as investigate the effects of specific features, such as past returns, on the hazard rate for selling the stock. Agnew, Balduzzi, and Sunden (2003) consider asset allocations and frequency of trading in investor 401(k) plans, however they focus more on broad statistics such as cross-sectional averages and their relation to investor characteristics, as well as correlations of allocation changes with past performance.

Our study will build on the existing research in several regards. The general goal is to produce a more in-depth analysis of how investor characteristics, market performance, and past portfolio performance give rise to certain features, which in turn influence the trading of an investor that ultimately affects the profit-and-losst dynamics of their portfolios. We hope to track the path dependence of an investor's experience over time and across a variety of environments so as to develop a more complete perspective of investor behavior in practically relevant contexts. The large size of the dataset, the longer time horizon it spans, and the increased computing power available today will enable us to effectively apply machine-learning approaches where this was not possible in earlier studies. We believe that these approaches will provide improved predictive ability over regressions, while at the same time allowing to identify relevant features in a more systematic fashion.

4.2 Data Description

We go into more detail about the data we are using. As mentioned before, it comes from a U.S. retail brokerage firm. Besides data on trades and positions, we also have records of flows into the accounts and investor demographics. The broker initially sent us all trading records for 2015 and only one snapshot of positions corresponding to the end of December 2015. While we will analyze the full 13-year dataset later in the project, the preliminary results pertain only to the 2015 dataset.

Table 4.1 summarizes the trading activity of all investors that are included in the positions snapshot. There are approximately 566,000 investors in total, carrying out over 6 million trades. A lot of these investors are inactive; only 45% of all investors made at least one trade in 2015. This lack of trading activity could be explained by investors holding passive investments, as well as by those who have a small balance of funds remaining with the broker and no longer use their account (while still having it open). For the investors who did trade at least once in 2015, we report the mean, the median, and percentiles for the distribution of the number of trades and the average trade. We again see that a large part of these investors are passive (40% of them traded 4 times or less in 2015), and at the same time there is presence of very active investors – those in the 95th percentile traded 83 times or more in 2015. There is also a wide dispersion in the size of trades carried out. The average trade for the median investor is \$5,000, while

Table 4.1: Summary of Investor Trading in 2015

# Investors (Dec 2015)			565,914
# Trades			6,357,800
# Trades per Investor			11.2
% Investors who traded at least once in 2015			45%
Cond. on at least one trade in 2015:			
Across Investors:		Number of Trades	Average Trade
Mean		24.9	$ 13,568
Median		6	$ 5,000
Percentile	1%	1	$ 49
	5%	1	$ 256
	10%	1	$ 535
	20%	2	$ 1,247
	40%	4	$ 3,381
	60%	9	$ 7,044
	80%	29	$ 15,957
	90%	55	$ 29,034
	95%	83	$ 48,603
	99%	261	$ 132,602

This table summarizes the trading activity of all investors for whom we have positions on record at the end of December 2015. Besides providing statistics for the whole dataset, we also look at the subset of investors who traded at least once in 2015, and summarize the distribution of the number of trades and average trade per investor for the whole year.

average trade for investors in the 95th percentile is almost $50,000. This dispersion in trading frequencies and sizes suggests that there are different classes of investors. The identification of these classes is beyond the scope of this thesis but is definitely something we plan to carry out later in our analysis.

An important initial part of the project has involved building the infrastructure to transform the raw data into a pre-processed format at which point the calculation of relevant features and subsequent analytics become much easier. More specifically, for each investor we want to know the number of shares held and the price of each security in their portfolio *on each trading day*. Our infrastructure is able to calculate these numbers by taking the portfolio positions snapshot and iteratively “rolling back" these positions to obtain the portfolio snapshot on the previous trading day. The roll-back procedure involves using the information on trades (which we have from the broker) and on historical prices.

CRSP is currently the only data source for prices that we are using. This obviously restricts us to how much of the investor’s portfolio we can price, considering that there are significant holdings and trading activity in other asset classes, such as fixed income and mutual funds. However, we can still obtain insights *at the portfolio level* by looking at just the equity holdings in an investor's portfolio.[1]

We restrict our analysis only to those investors for whom we can price *all* their equity holdings in the December 2015 portfolio snapshot. This filter still leaves us with a large dataset of 256,986 investors and 7,632 distinct securities for which we have prices available in CRSP. We note that there is an implicit survivorship bias associated with this procedure whereby we use end of year data and filters to look at trading activity earlier in the year. Unfortunately, we cannot do better since we only have the portfolio snapshot for December 2015. That being said, the bias should not be too big since investors rarely close their accounts, while the CRSP universe captures most of the equity securities (and almost all U.S. common stocks).

In the rest of the paper we will refer to a *portfolio* as the collection of all equity securities held by the investor; we do not include cash here.

4.3 Trading Distribution Analysis Methodology

We are interested in analyzing how past portfolio and security returns relate to subsequent trading. This is done by calculating the distribution of the *net trade* by an investor over a specified future horizon f conditional on *past returns* measured over a specified past horizon b. We also condition on if a trade took place over the past horizon or not. We next outline how the past returns, the net trade, and the trading indicator are computed.

4.3.1 Computing Past Returns

There are two popular ways for calculating portfolio return: the *time-weighted return* approach and the *modified Dietz* approach.

For calculating the *time-weighted return* we look at the change in portfolio balance in each period excluding cash flows, and compound these changes over time. Denote by B_t the portfolio balance at the end of day t and by CF_t is the cash flow attributed to day t.

[1]Each trade and position in the dataset includes an asset group classification for the security associated with that trade/position. This makes it easy to identify all equity holdings and trades.

The return over a particular horizon, e.g. from day $t-b$ to day t, is deifned as:

$$r_t = \prod_{s=t-b+1}^{t} \frac{B_s - CF_s}{B_{s-1}} - 1 \tag{1}$$

The advantage of this formula is that is gives an accurate measure for the return on the assets held in the portfolio. The disadvantage is that is does not account for the fact that some daily returns may contribute more to the overall return than others. For example, consider the following two cases. In both cases an investor holds one share of stock A worth \$10 on day 0, which returns 0% on day 1 and 100% on day 2. In the first case the investor does not make any trades, so their portfolio at the end of day 2 is worth \$20. In the second case the investor buys four more shares of A at the end of day 1, so their portfolio at the end of day 2 is worth \$100. Even though the time-weighted return formula gives the same number in both cases, it is clear that the investor has a significantly greater portfolio return in the second case.

The *modified Dietz* approach addresses the above short-coming as follows. It aggregates all cash flows at once, giving a higher weight to the ones that have been held in the portfolio for longer. The formula is:

$$r_t = \frac{B_t - B_{t-b} - \sum_{s=t-b+1}^{t} CF_s}{B_{t-b} + \sum_{s=t-b+1}^{t} w_s CF_s} \tag{2}$$

where $w_s = (t-s)/b$ is the weight corresponding to the cash flow on day s. It is important to note that this formula is an approximation since it does not look at the incremental changes in portfolio level (which is done for the time-weighted return).

Which formula should be used? This depends on the application at hand. Here, we are dealing with the equity portion of an retail investor's portfolio, where it is not uncommon to see large cash flows (as we will show later). Furthermore, it is important to understand how the investors themselves go about computing their portfolio returns, since this is what ultimately drives their decision-making. Again, since this is a retail group, it is reasonable to assume that most investors just focus on the *levels* of equity balances rather than the granular cash flows adjustments that appear in the time-weighted formula. Based on this intuition we employ the modified Dietz method for computing past returns.

We go into more detail about how formula (2) is computed in our setting. The balance B_t is just the dollar value of an investor's *equity portfolio* at the close of day t, which is the total dollar amount of all equity security positions in their portfolio. Cash flows are calculated by accounting for all trades made over the period.[2] More specifically, if we denote by $q_{j,s}$ the number of shares held in security j at the end of day s, and by $p_{j,s}$ its closing price, then the cash flow on day s is:

$$CF_s = \sum_{j=1}^{N} (q_{j,s} - q_{j,s-1}) p_{j,s} \tag{3}$$

where the securities in the equity portfolio are indexed from 1 to N. This way, the cash flow is equal to the aggregate dollar trade across all securities on a particular day. For security level analysis, the above formula still applies. For calculating the return for a particular security j, we use for the balance B_t just the position in that security, calculated as $q_{j,t} \times p_{j,t}$. For the cash flows calculation we use the trades in that security only.

4.3.2 Computing Net Trade

We next turn to quantifying trading behavior. Let us denote by a *net trade* as the shift in an investor's portfolio as a result of trading. In particular, a low (and negative) number should correspond to an investor selling their holdings, whereas a higher (and positive) number should correspond to adding to existing positions and/or buying new securities.

At the security level, this measure is very intuitive. We restrict our analysis to cases in which an investor is holding a positive number of shares in a security on day t. We then look at the percent change in the number of shares over the future horizon:

$$d_t = \frac{q_{j,t+f} - q_{j,t}}{q_{j,t}} \tag{4}$$

Note that if this formula gives a value of 0%, this corresponds to no change in the position,

[2] We do not need to consider dividends with this formula. This is because if a dividend is not reinvested, it would move into the cash portion of the portfolio and not affect the equity portfolio balance, and if it is reinvested, then the number of shares held would be adjusted and would then be reflected directly in the portfolio balance. Also, dividends will generally only marginally affect the whole investor portfolio balance.

whereas as value of -100% corresponds to a liquidation of a position.

At the portfolio level things are more complicated since there are multiple securities with potentially substantially different prices. There are two approaches that we consider. The first generalizes the security level formula (4) by weighting the percent changes in holdings of securities by their current balance in the portfolio:

$$d_t = \sum_{j=1}^{N} w_j \frac{q_{j,t+f} - q_{j,t}}{q_{j,t}} \tag{5}$$

where $w_j = (q_{j,t} \times p_{j,t})/B_t$ is the security weight in the portfolio at day t. We need to slightly modify this formula to account for securities that are not held in the portfolio on the current day but traded sometime over the future horizon. The modification involves weighting the *share changes* in holdings of securities by their current *prices*:

$$d_t = \frac{\sum_{j=1}^{N} (q_{j,t+f} - q_{j,t}) p_{j,t}}{B_t} \tag{6}$$

where we now sum over all securities that are held in the portfolio on at least one day from the current day to the future horizon, and not just the ones held on the current day.[3] Again, it is not hard to show that the net trade in (6) is equal to -100% if and only if an investor completely liquidates their portfolio.

We also propose an alternative approach which aggregates trading at the dollar level. This measure is useful since it explicitly links cash flows to their effect on the portfolio. It is important to adjust for security returns when aggregating the cash flows. For example, consider the following situation. An investor is holding one share of stock A worth \$10 on day 0, which returns 100% on day 1. At the end of day 1 the investor completely liquidates their position, corresponding to a trade of $-\$20$. This trade is equal to -200% of the day 0 portfolio balance, even though the investor traded only 100% of their portfolio.

To make the cash flows on different days comparable to each other, we *discount* all cash flows by the time-weighted portfolio return. We then add the discounted cash flows and divide by the portfolio balance to give a measure of net trade. The formula is:

$$d_t = \frac{\sum_{s=t+1}^{t+f} v_s CF_s}{B_t} \tag{7}$$

[3]If no price if available for a security on day t (e.g. due to an IPO that occurs later in the sample), we assume this price is equal to the first price in the future that is available.

where:

$$v_s = \Big[\prod_{u=t+1}^{s} \frac{B_u - CF_u}{B_{u-1}} \Big]^{-1} \tag{8}$$

This formula is equivalent to the shared-based formula (6) if we discount the cash flows for each security by the corresponding time-weighted returns on the security position. Which of the two discounting approaches to use comes down to the prior on investor behavior – i.e. when an investor decides on a trade in *one security*, do they care about the portfolio level or just the position in that security. We will investigate the performance of both methods in the next section.

4.3.3 Computing Trading Indicator

Past trading behavior is another important feature that should be considered when predicting future trading. Barberis and Xiong (2012) propose the idea of realization utility, whereby investors experience shocks to their utility after *realizing* losses and gains due to trading. It may also be the case that following a trade in a security, investors start paying more attention to it (at least in the short-term).

We quantify past trading in a very simple fashion by computing a *trading indicator*, which takes the value of 1 if an investor traded at least once over the past horizon, and a value of 0 otherwise. Of course, the case in which a person made at least one trade may involve very distinct scenarios, such as exactly one large purchase of a security, several trades in the same direction corresponding to building of a position, or regular frequent trading by an active investor. Stratifying past trading history into more cases is something we plan to do later in the project.

4.3.4 Computing Conditional Distributions

Finally, we discuss how the conditional distributions are computed. Out of our filtered sample of 256,986 investors we exclude the ones who never traded in 2015. A lot of these are likely inactive accounts, and while we introduce a further look-ahead bias with this filter, in the end we are interested in characterizing the behavior of people who do trade. If we include inactive accounts in the sample, then we will be systematically underestimating the conditional probabilities of trading. We have checked that the empirical results and conclusions do not change significantly depending on whether or not we enforce this filter.

Distributions across investors are computed by assigning past returns and net trades to disjoint intervals. We discuss these intervals in the Appendix.

At the security level, for each of the 132,665 investors in the resulting sample, we look at all securities that they traded in 2015 and that we can price in CRSP. For every such investor-security pair, we look at each trading day t and compute the past return and the net trade over the future horizon for all observations where the number of shares held is positive. We also compute the trading indicator as outlined before. We then count the observations corresponding to:

- Each past horizon interval
- Each future net trade interval
- Each trading indicator

Based on the observation counts the conditional distributions are straightforward to compute. For portfolio level analysis the procedure is similar. The only difference is that we impose a further filter whereby among the 132,665 investors from before we only consider those for whom we can also price all the equity securities they traded in 2015. The filter results in 125,724 investors, still a very large sample. After that we compute the portfolio past return and net trade on each trading day t for each investor. The trading indicator is also now measured at the portfolio level (taking the value 1 if the person traded at least once in the equity portfolio).

4.4 Trading Distribution

Having outlined our methodology for quantifying trading behavior, we are ready to explore the distribution of net trades. We start by looking at the unconditional distribution, both at the security and at the portfolo level. We find that trades in individual securities predominantly comprise of position liquidations, whereas trades in portfolios are more clustered around 0% of the portfolio notional. We provide potential explanations on how to reconcile this difference in distributions.

Figure 4-1 plots the distribution of net trades at the security level *assuming* that a non-zero net trade occurred. We see that liquidating a position is a very common trade; it corresponds to just over 35% of observations for a 1-day horizon and almost 45% of observations for a 60-day horizon.[4] We also observe a high probability of net trades being around -50% and around 50%, as well as in the right tail around $100\%, 150\%, 200\%, 300\%$, and 400%. This is driven by two factors. The first is trading in round lots, which is quite well-documented in the literature; see, for example, Moulton (2005) and Alexander and Peterson (2007). If an investor is currently holding 100 shares in a stock, and they trade a round number of shares, then any change in their position would be a "simple" fraction, such as 20% if they buy 20 shares or -40% if they sell 40 shares. We believe another important factor is the use of behavioral heuristics by investors. An investor who decides to put on a trade will often not calculate the optimal change in his position, but rather choose some round intuitive number representing a *proportion of their position*, such as one half or one third, or some integer multiple.

In Figure 4-1 we examine the right tail of the distribution, for trades greater than 100%. The total mass of the distribution in this tail is substantial, making up around 5% of all observations, both at the 1-day and the 60-day horizons. Furthermore, investors increase their positions by more than 500% about 0.7% of the time for the 1-day horizon and 0.8% of the time for the 60-day horizon. This again is a significant proportion of observations.

While we calculate the distribution over all trades in the sample, one may think that the distribution is skewed by the population of active investors since they carry out more trades. For robustness, we have repeated the exercise for investors who traded less than 10 times in 2015 and separately for investors whose average number of trades per security is less than 3. In both cases the results don't materially change.

We next move to the portfolio level analysis. Figure 4-2 plots the trade distribution when we use the share-based measure in (6) for calculating the net trade. The distribution shows a very heavy clustering of mass around a 0% change in the portfolio and a monotone decrease in the likelihood as we move away from 0%. This distribution is at a stark contrast with what we saw for individual securities, suggesting that the decision-making process is very different depending on if an investor is thinking how much to trade in one security versus how much to trade as a proportion of their entire portfolio. Another reason why we are getting a different distribution is that at the portfolio level we are now accounting for buying of new securities.[5]

[4]We observe the same result when considering horizons at 5, 10, 15, ..., 55 days.

[5]At the security level we assume the current position is positive, so we only restrict to securities

The tails are also important in the portfolio net trade distribution. From Figure 4-2 we see that portfolio liquidation is still apparent, although to a much smaller extent than for securities. For a 1-day horizon, liquidation occurs 1.8% of the time; this drops to 1.2% for a 60-day horizon. The right tail is thicker than for the security level distribution; at the 60-day horizon there are over 6% of observations corresponding to net trades greater than 100%, and almost 1.4% of observations corresponding to net trades greater than 500%. This increased right tail is likely caused by purchases of new securities.

We also examine how the dollar-based net trade measure performs at the portfolio level; Figure 4-7 in the Appendix plots its distribution. The distribution looks very similar to what we saw with the share-based approach, both in terms of the shape and in terms of magnitudes. There is a slight difference in the left tail, whereby we now see a few trades that are less than −100% of the portfolio. (Recall that with the share-based method these trades never occur in absence of shorting). The right tail is substantially different and much larger; at the 60-day horizon over 7.5% of net traders are greater than 100%, and almost 3.6% are greater than 500%. This is likely due to the fact that discounting by the time-weighted portfolio return overweights new purchases in securities that significantly rise in price over the future horizon – in comparison to the share-weighted approach in which such purchases are discounted by the security price return.

Overall, it is fair to conclude that the share-based and the dollar-based measures are similar to each other, and which one to use depends on how we want to model the tail behavior of net trades. For the remainder of this paper we will always employ the share-based measure, in part because it provides a more consistent comparison between the portfolio and the security level trading.

We discuss some reasons for why there is such a difference between the trading distributions at the portfolio and at the security levels. The first is the existence of "reshuffling" in investor portfolios, whereby investors will at times choose to close out positions in securities they no longer wish to hold and replace them with other securities instead. This is related to the rank effect documented by Hartzmark (2015), whereby investors tend to predominantly sell extreme performers in their portfolios. Another reason could be the presence of managed accounts in the sample, which tend to systematically rebalance in a similar fashion as to what an institutional investor would do. Furthermore, a large part

already held in the investor's portfolio.

of the trades have a fixed dollar commission associated with them, which may encourage investors to trade in large quantities rather than to gradually adjust their portfolios. We plan to explore these aspects in more detail later on, especially when we get the full dataset.

To provide further validation that the two distributions are different, we split the net trade distribution at the security level into two cases: when an investor traded exactly one security in 2015 and when they traded at least two securities. Figure 4-3 plots the resulting distributions, both at the 1-day and the 60-day future horizons. It is evident that the net trade distribution for the case when at least two securities were traded looks very similar to the overall security distribution in Figure 4-1. In contrast, when an investor traded exactly one security, the distribution of their trades resembles the portfolio level distribution; this is especially apparent as we get to the longer horizon of 60 days. This suggests that an investor with just one security in their portfolio views their trading as occurring at the portfolio level (even though they are trading just one security). This may be rational behavior if the security they are holding is diversified in itself, such as a broad market ETF. At the same time, if an investor holds 20 different single name stocks, it may make sense for them to liquidate one of these stocks and replace it with another when they wish to rebalance.

To summarize, our insights into the overall trading distribution are as follows. A very large proportion (as much as 45%) of trades in individual securities are position liquidations. We also observe trading at round fractions and multiples of positions at the security level, which may in part be explained by behavioral investor heuristics. When it comes to portfolio level trading, most trades are close to 0% of portfolio notional, although we still see a non-trivial amount of liquidations and substantial increases in portfolio positions. We may attribute the stark difference in the portfolio and security level trading distributions to the fact that investors liquidate positions in some securities and replace them with other securities. Further work is needed to understand this behavior and to investigate other reasons for the difference.

Figure 4-1: Net Trade Distribution at Security Level

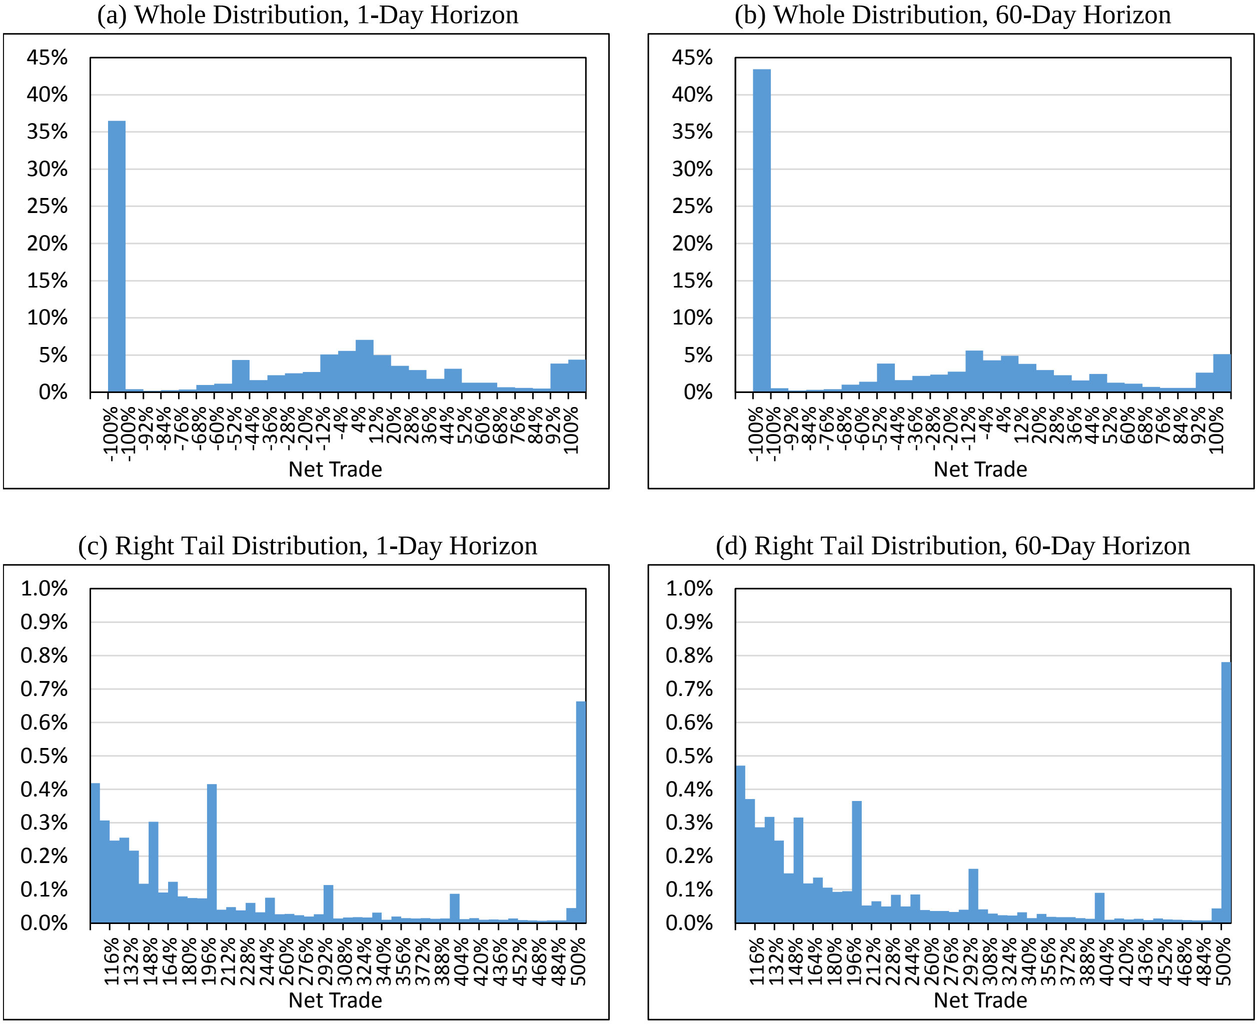

This figure plots the unconditional distribution of the net trades over the future horizon. Analysis is done at the *security* level. We consider the 1-day and the 60-day future horizon. The net trade is measured as the change in number of shares held in the security, divided by current number of shares.

The top two charts show the distribution over the whole interval $(-\infty, \infty)$, whereas the bottom two show the right tail distribution over the interval $(100\%, \infty)$. For each chart the distribution is computed by calculating the probability of the net trade falling into a particular interval. The endpoints of the intervals are shown on the $x-$axis.

All intervals have length of 8% except the following. For the top charts, the left-most interval corresponds to a net trade of less than -100%, the second left-most interval corresponds to a net trade exactly equal to -100%, while the right-most interval corresponds to a net trade of greater than 100%. For the bottom charts the right-most interval corresponds to a net trade of greater than 500%. Standard errors for the empirical probability corresponding to each interval are at most 0.02% (except the 1-day liquidation probability at 0.04%) and are not shown in the charts.

Figure 4-2: Net Trade Distribution at Portfolio Level
Share-Based Approach

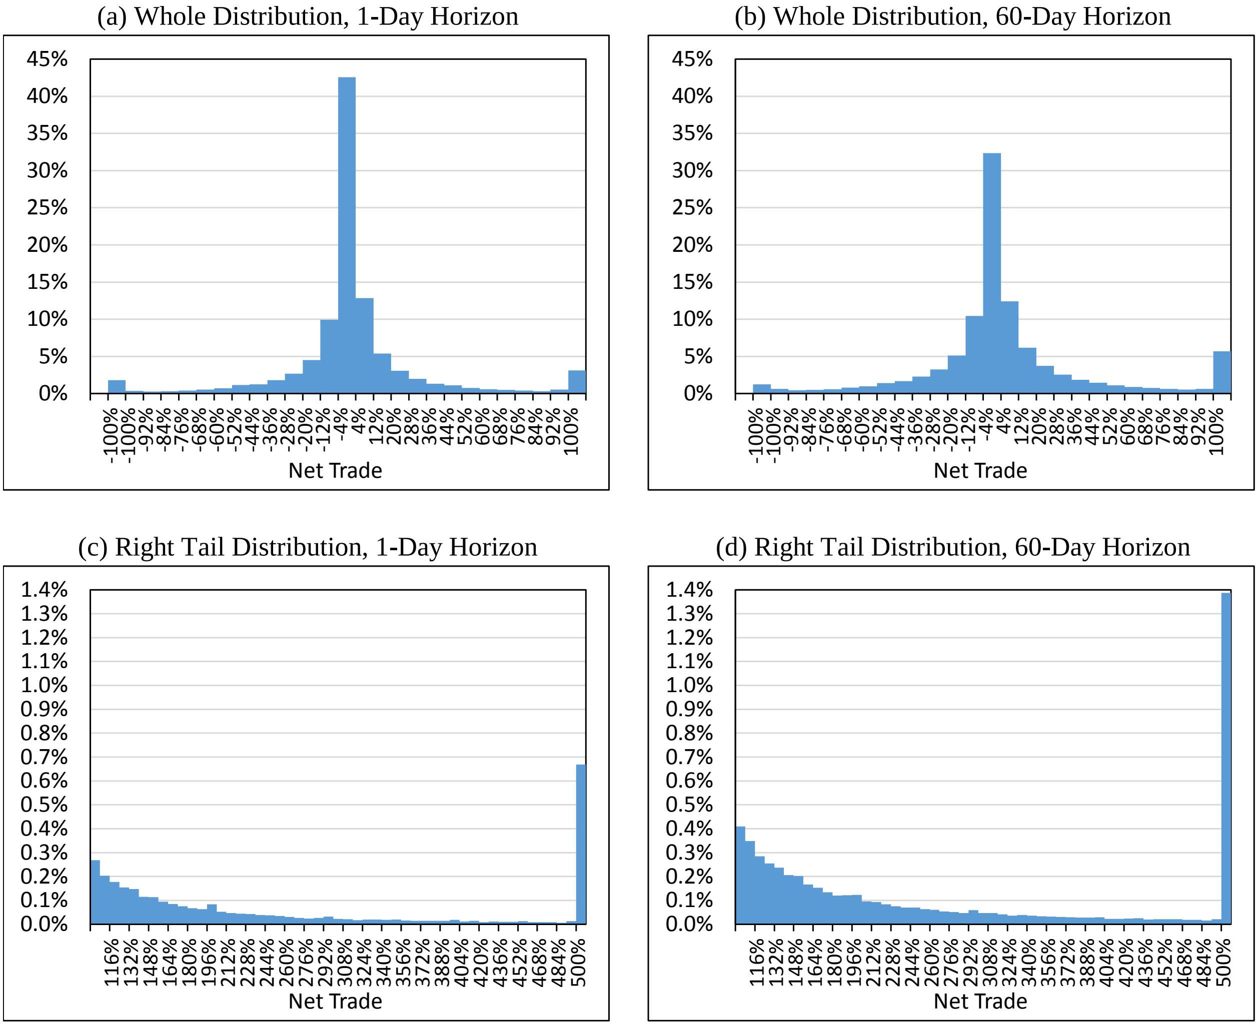

This figure plots the unconditional distribution of the net trade over the future horizon. Analysis is done at the *portfolio* level. We consider the 1-day and the 60-day future horizon. The net trade is calculated as an average change in number of shares held in each security, weighted by its current price, as in equation (6).

The top two charts show the distribution over the whole interval $(-\infty, \infty)$, whereas the bottom two show the right tail distribution over the interval $(100\%, \infty)$. For each chart the distribution is computed by calculating the probability of the net trade falling into a particular interval. The endpoints of the intervals are shown on the $x-$axis.

All intervals have length of 8% except the following. For the top charts, the left-most interval corresponds to a net trade of less than -100%, the second left-most interval corresponds to a net trade exactly equal to -100%, while the right-most interval corresponds to a net trade of greater than 100%. For the bottom charts the right-most interval corresponds to a net trade of greater than 500%. Standard errors for the empirical probability corresponding to each interval are at most 0.02% (except for the intervals between -12% and 12%, where the bound is 0.05%) and are not shown in the charts.

Figure 4-3: Comparison of Net Trade Distribution at Security Level Depending on Number of Securities Traded in 2015

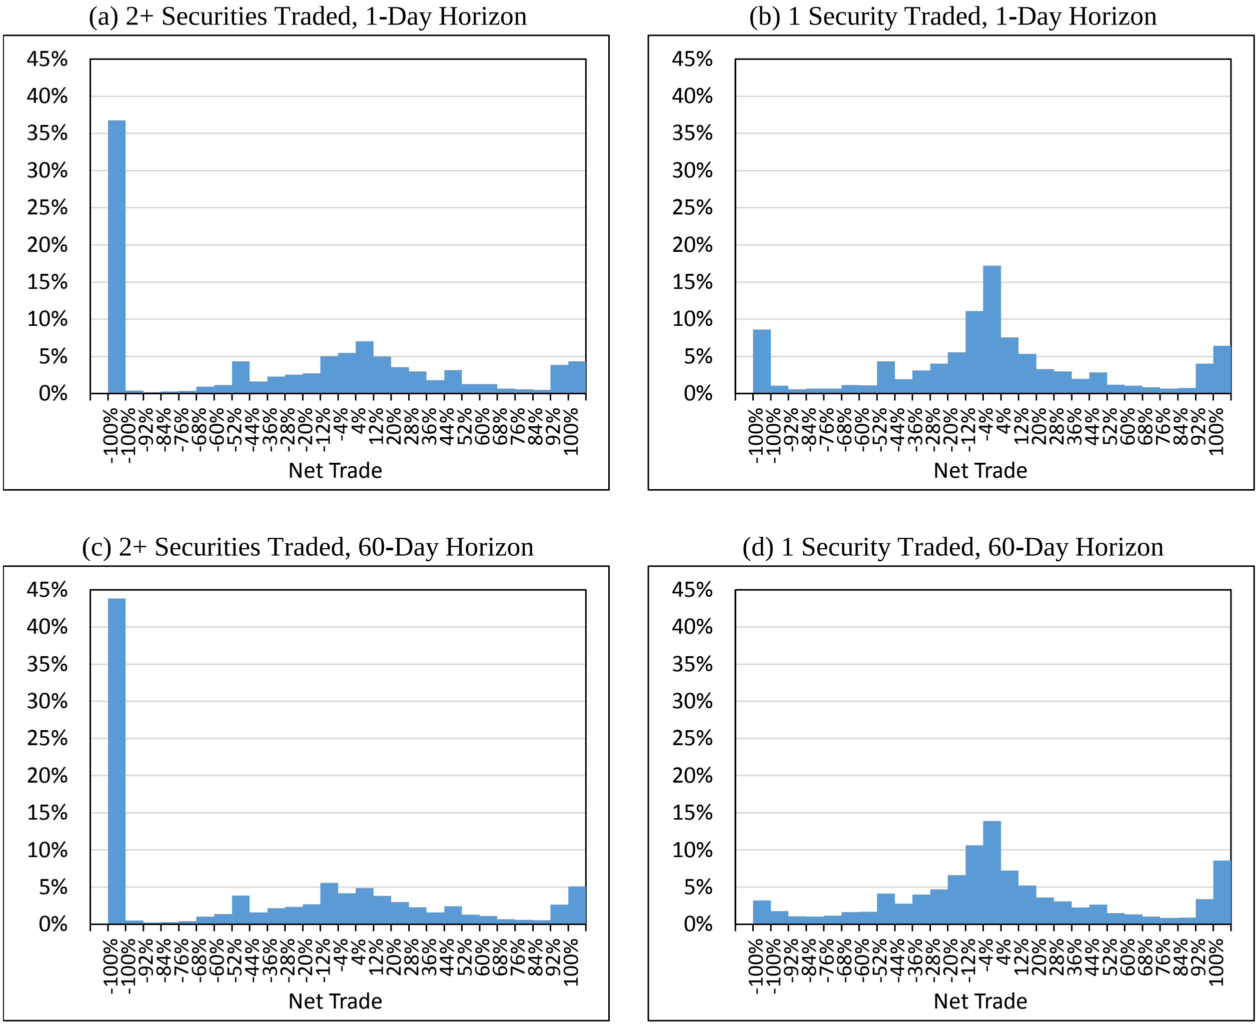

This figure plots the trading distribution at the *security* level for two scenarios: when an investor traded at least two securities in 2015 and when they traded exactly one security. We consider the 1-day and the 60-day future horizon. The net trade is measured as the change in number of shares held in the security, divided by current number of shares.

For each chart the distribution is computed by calculating the probability of the net trade falling into a particular interval. The endpoints of the intervals are shown on the $x-$axis. All intervals have length of 8% except the following. For the top charts, the left-most interval corresponds to a net trade of less than -100%, the second left-most interval corresponds to a net trade exactly equal to -100%, while the right-most interval corresponds to a net trade of greater than 100%. For the bottom charts the right-most interval corresponds to a net trade of greater than 500%. Standard errors for the empirical probability corresponding to each interval are at most 0.32% and are not shown in the charts.

4.5 Dependence on Past History

In order to build predictive models of investor behavior, we need to decide on what features to include in the models. Here, we focus on two important features – past returns and past trading behavior – and analyze the distribution of net trades *conditional* on these two attributes. The key result is that the probabilities of liquidation and of significant increases in positions are much larger following extreme returns; this is consistent with earlier work by Ben-David and Hirshleifer (2012). We also investigate how these reactions to large losses and large gains depend on whether or not an investor traded in the past, as well as what horizons are used for measuring past returns and subsequent trading.

We start by stratifying the distribution of net trades by the returns over the past horizon; we use 120 days for the past horizon and 60 days for the future horizon. Table 4.2 shows the resulting conditional distributions at the security level. Overall, the distributions all look similar to the unconditional distribution we discussed earlier in Table 4-1, namely that there is a large concentration of trades corresponding to position liquidations, as well as some concentration around 0% and in the right tail. However, by comparing the different rows in the table, we see that following more extreme returns, liquidations make up a much larger proportion of all trades: 57% following a loss of -39% or more and 52% following a gain of 39% or more, in comparison with just 35% following past returns in the range $(-3\%, 3\%]$. It is also evident that *large buys*, which we define as increases in position by 100% or more, make up a larger proportion of all trades after large losses or large gains. At the same time, the proportion of trades that are close to 0% declines substantially.

We repeat the same exercise at the portfolio level; Table 4.3 presents the results. While the distributions are now more clustered around 0% (a phenomenon we observed earlier when looking at the unconditional distribution in Table 4-2), we again see the distribution mass shifting to the tails following returns that are larger in magnitude. For example, if past returns are in the range $(-3\%, 3\%]$, then 59% of non-zero trades fall into the $(-10\%, 10\%)$ interval. This number drops to just 18% if past returns are less than -39% and to 21% if past returns are greater than 39%. Consistent with this behavior, we also see that the proportion of trades corresponding to portfolio liquidations and to large buys increases after extreme returns.

To verify that this result holds for other choices of past and future horizons, we repeat the same exercise for four other horizon cases and plot the results in three-dimensional charts in Figures 4-8 and 4-9 in the Appendix. We see that the dependence of the distributions on past returns is very similar (and holds even if we split into two separate cases for whether a trade was made over the past horizon).

So far we have restricted our observations to those where the net trade is non-zero. We now look at *all* observations following a particular past return and ask the question of how many of those observations correspond to a certain trade. This is the more important question from a predictive perspective, since we want to forecast if an investor will make a trade (or a trade of a certain kind) based on past history. For this paper we will focus on two trades: a net trade of -100% corresponding to a *liquidation*, and a net trade of 100% or more corresponding to a *large buy*.[6]

Figure 4-4 plots the probability of security liquidation conditional on past returns. This probability is much higher following large losses and large gains. Following losses of over 39% over the past 20 days, the conditional probability of liquidation over the next 60 days is 19%, in comparison to just 7% following past returns between -3% and 3%. The conditional probability following gains of 39% or more is also very high, at 22%. A conditional probability number around 20% is actually impressive in magnitude, because it means that by conditioning just on one feature of past data, we can predict with 20% probability that a person will make a very specific and an important trade over the future horizon.

For robustness, we also plot the conditional probability if we measure net trade over the next 5 days instead of the next 60 days; we again see more extreme reactions following large returns. The same phenomenon is observed for probabilities of large buying, also shown in Figure 4-4; the overall levels are lower than for liquidation probabilities, since buying additional shares of a security is quite rare. We thus see a very consistent "V-shaped" curve for selling and buying probabilities conditional on past returns. Ben-David and Hirshleifer (2012) document this result and note that it lies in contrast with predictions due to the disposition effect, whereby investors are reluctant to realize losses and hold on to losing investments for too long while selling winning ones too soon.

The result also holds if we look at the portfolio level conditional probabilities; see

[6] While these two trades are important, there are other types one may be interested in, especially at the portfolio level – such as increases in positions by 10% or more or decreases by 10% or more. This is beyond the scope of the paper.

Figure 4-10 in the Appendix. The probabilities of large buys are now higher since for portfolios we account for buying of new securities.

We next investigate the dependence of net trades on another historical feature, which is past investor trading. In Figure 4-5 we plot the probability of liquidation conditional on past returns for two cases: if at least one trade was made over the past horizon and if there were no trades. We see that if an investor traded in the past, they are a lot more likely to liquidate their position. In more detail, for a past horizon of 20 days and future horizon of 60 days, the probability of liquidation after a loss of -39% or more is 31% if a trade was made, and 11% is no trade was made in the past. A similar picture if observed if we consider the shorter future horizon of 5 days; the correspoding probabilities are 8% if a trade was made in the past, and just 1% if no trade was made. We obtain the same result if we compare the probabilities of large buys for the two cases in Figure 4-5.

In Figure 4-6 we repeat the same procedure at the portfolio level. Overall, the result still holds, with past trading leading to a higher probability of trading in the future. We next compare the probabilities of liquidation conditional on past returns for securities and for portfolios. The levels of the probabilities are higher for securities, which is consistent with the unconditional distribution comparison from before. However, we see an interesting phenomenon whereby the liquidation probability after large losses *relative* to flat returns is greater at the portfolio than at the security level, where we define *large losses* as past returns less than or equal to -39% and *flat returns* as past returns falling in the $(-3\%, 3\%]$ interval. For example, take the case when an investor traded in the past, the past horizon is 20 days, and the future horizon is 60 days. (The results for other cases are similar). From Figure 4-5(a) we see that the probability of security liquidation following large losses is 31%, a three-fold increase over the probability of 11% following flat returns. In contrast, from Figure 4-6(a) for portfolios, the liquidation probability after large losses is 4.4%, which is over 8 times higher than this probability after flat returns of just 0.5%. Furthermore, this means that in “normal conditions" of past portfolio returns that are flat, investors very rarely liquidate their whole portfolios, whereas when their portfolio returns are extreme, the liquidation probability becomes substantial.

We can further extend the comparison of probabilities to other situations, such as looking at buying vs. selling behavior or looking at reactions to losses vs. gains. We do not yet have a conclusive approach to this question but want to provide our initial thoughts here. There are a few different quantitative measures one could use for the comparison, such as the average net trade following a particular return, or the dispersion of net trades. Another idea is to "summarize" the shape of the conditional probability "V-curve", for example by looking at if the curve is higher for large gains or large losses, or by "scaling" the curve so that it corresponds to a probability distribution, and then comparing this distribution to a uniform one (because one may set the null hypothesis as an investor's decision being independent of past returns).

We go into more detail about one way to summarize the V-shape, which is to measure its steepness. We can formally define the *steepness* of the liquidation probability curve on the losses side as the conditional probability of liquidation following large losses divided by the conditional probability of liquidation following flat returns. We then analyze how this steepness varies for various choices of past and future horizons and present the results in Figure 4.4 at both the security and the portfolio level. We see that the V-shapes are steeper for shorter horizons and flatter for longer horizons; the steepness is also more pronounced for portfolios than for securities, which is what we saw earlier. Of course, more work needs to be done to investigate the effects of returns *at different horizons* on trading.

**Table 4.2: Net Trade Distribution Conditional on Past Security Returns
Past Horizon 120 Days, Future Horizon 60 Days**

Conditional Distribution in rows	Net Security Trade over Next 60 Days																				
Return over Past 120 Days	-100%	-90%	-80%	-70%	-60%	-50%	-40%	-30%	-20%	-10%	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
-39%	57.18	0.43	0.36	0.36	0.66	1.75	3.53	1.15	1.26	1.03	0.90	2.09	3.53	2.46	2.72	1.74	2.90	1.91	0.95	0.85	0.84	11.40
-33%	55.97	0.60	0.42	0.40	0.81	1.95	1.40	1.33	1.43	1.23	1.22	3.26	4.53	3.24	3.36	1.91	2.93	1.93	0.97	0.80	0.96	9.33
-27%	53.95	0.70	0.38	0.45	0.77	1.95	2.53	1.63	1.76	1.56	1.62	3.64	4.98	3.59	3.56	1.85	2.53	1.87	1.10	0.74	0.85	7.99
-21%	53.88	0.69	0.33	0.35	0.68	2.01	2.11	1.70	2.20	1.78	1.94	4.12	5.93	3.88	3.38	1.63	2.31	1.59	0.95	0.64	0.76	7.15
-15%	52.63	0.65	0.31	0.44	0.74	2.16	3.47	1.65	2.23	1.95	2.07	5.29	6.64	3.88	3.21	1.49	2.00	1.40	0.70	0.62	0.72	5.73
-9%	46.39	0.64	0.37	0.51	0.77	2.06	4.26	1.93	2.49	3.14	4.24	7.66	6.61	3.71	2.99	1.51	1.89	1.24	0.66	0.61	0.70	5.63
-3%	40.63	0.70	0.34	0.49	0.82	2.23	3.43	2.17	2.81	3.90	6.59	8.98	5.84	3.61	2.57	2.29	1.88	1.12	0.61	0.66	0.98	7.34
3%	35.23	0.64	0.33	0.44	1.12	2.77	3.89	2.15	2.58	3.71	11.41	6.02	4.23	3.66	2.20	4.46	5.66	1.51	0.56	0.59	0.87	5.97
9%	38.11	0.67	0.47	0.56	1.03	2.90	4.41	3.29	4.04	4.95	7.24	6.42	5.55	3.50	2.88	1.84	2.01	1.44	0.78	0.76	0.83	6.32
15%	39.98	0.68	0.51	0.62	1.11	3.17	2.59	3.25	4.25	5.26	5.61	5.05	5.36	3.83	3.27	1.89	2.33	1.53	0.84	0.77	0.91	7.20
21%	41.11	0.64	0.52	0.64	1.10	3.42	2.82	3.00	4.39	5.14	3.40	4.10	5.32	3.85	3.41	1.87	2.74	1.77	0.90	0.80	0.95	8.11
27%	42.15	0.58	0.48	0.62	1.12	3.68	3.74	2.84	4.16	4.08	2.80	3.54	5.38	3.41	3.31	1.93	2.98	1.71	0.92	0.85	0.96	8.79
33%	44.12	0.65	0.52	0.65	1.12	3.21	3.20	2.71	3.11	3.65	2.58	2.89	5.38	3.48	3.67	1.69	3.03	1.71	0.87	0.92	0.92	9.93
39%	47.00	0.68	0.56	0.67	1.25	3.12	2.89	2.64	2.84	3.17	2.78	2.32	3.92	3.13	3.42	1.68	3.08	1.76	0.94	0.82	0.85	10.49
	52.13	0.50	0.42	0.53	0.88	2.35	1.46	1.54	1.96	2.17	1.78	2.22	3.52	2.62	2.63	1.93	3.02	1.99	0.96	1.00	1.04	13.34

This table lists the distribution of the net trade over the future horizon conditional on past returns. Analysis is done at the *security* level. We assume the long horizon scenario, with a past horizon of 120 days and a future horizon of 60 days.
Each row corresponds to a conditional distribution of trades following the return over the past horizon falling into a particular interval. The partition points for these intervals are shown in the second column and range from −39% to 39%. The distribution itself is computed by calculating the conditional probability of the net trade falling into a particular interval. The partition points for the net trade intervals are shown in the second row and range from −100% to 100%. The cells contain the conditional probabilities as percentages without the percentage sign. We provide the color coding of the cells to the right of this caption.

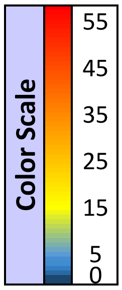

Table 4.3: Net Trade Distribution Conditional on Past Portfolio Returns Past Horizon 120 Days, Future Horizon 60 Days

Conditional Distribution in rows	Net Portfolio Trade over Next 60 Days																				
Return over Past 120 Days	-100%	-90%	-80%	-70%	-60%	-50%	-40%	-30%	-20%	-10%	0%	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
-39%	3.34	2.04	1.50	1.47	1.86	2.17	3.06	3.64	4.30	5.86	8.11	10.17	7.29	5.66	4.35	3.84	3.23	2.83	2.23	1.62	1.23	20.17
-33%	2.44	1.82	1.36	1.31	1.97	2.01	2.99	3.33	4.73	6.81	10.74	14.17	7.76	5.90	4.77	3.79	2.94	2.17	1.61	1.37	1.21	14.81
-27%	1.74	1.59	1.05	1.34	1.80	2.31	2.96	3.39	5.29	7.36	12.25	15.80	8.70	5.39	4.96	3.34	2.59	1.81	1.61	1.29	1.05	12.39
-21%	1.12	1.48	0.93	1.16	1.56	2.05	2.75	3.55	4.95	7.70	13.83	18.86	9.28	5.63	4.82	2.88	2.18	1.71	1.42	1.28	0.97	9.86
-15%	0.78	1.05	0.85	0.93	1.37	1.87	2.47	3.34	4.89	7.91	16.78	22.70	9.08	5.39	3.77	2.59	1.77	1.44	1.20	0.90	0.77	8.13
-9%	0.42	0.65	0.56	0.64	0.91	1.34	1.80	2.59	4.10	7.46	21.12	33.04	8.23	4.19	2.61	1.83	1.17	0.88	0.76	0.56	0.48	4.65
-3%	0.40	0.58	0.48	0.58	0.83	1.17	1.56	2.34	3.71	6.97	23.14	35.88	7.85	3.87	2.29	1.57	1.04	0.75	0.60	0.44	0.41	3.53
3%	0.53	0.62	0.54	0.53	0.84	1.20	1.58	2.20	3.45	6.69	29.63	29.83	7.68	3.85	2.40	1.52	1.05	0.74	0.60	0.47	0.43	3.62
9%	0.95	0.87	0.72	0.74	1.15	1.60	2.11	2.75	4.24	7.63	21.33	26.76	9.49	4.80	3.10	1.98	1.45	1.11	0.86	0.62	0.63	5.12
15%	1.28	1.14	0.85	1.02	1.28	2.05	2.45	3.27	4.61	7.53	16.66	24.01	9.41	5.27	3.59	2.43	1.89	1.41	1.15	0.72	0.71	7.26
21%	1.79	1.42	1.13	1.20	1.77	2.50	2.77	3.41	4.73	7.44	14.21	19.91	8.76	5.52	3.71	2.58	2.06	1.81	1.50	1.11	0.80	9.88
27%	2.58	1.36	1.21	1.30	2.08	2.83	3.11	4.11	4.48	6.86	12.97	16.61	7.87	5.92	3.61	2.87	2.16	1.86	1.61	1.24	1.07	12.31
33%	3.13	1.72	1.14	1.33	2.26	2.68	2.85	4.56	4.37	6.46	13.09	14.16	7.29	5.88	3.74	3.17	2.29	1.96	1.18	1.36	1.22	14.16
39%	3.71	1.90	1.17	2.16	2.01	3.27	2.95	4.09	4.64	6.87	13.07	12.68	6.87	5.40	3.16	2.75	1.71	2.23	1.77	1.41	1.29	14.90
	4.84	2.85	0.86	1.28	2.21	2.79	2.72	3.25	4.24	6.19	9.89	10.90	7.32	4.88	2.97	2.57	2.57	2.28	2.03	1.59	1.46	20.30

This table lists the distribution of the net trade over the future horizon conditional on past returns. Analysis is done at the *portfolio* level. We assume the long horizon scenario, with a past horizon of 120 days and a future horizon of 60 days.
Each row corresponds to a conditional distribution of trades following the return over the past horizon falling into a particular interval. The partition points for these intervals are shown in the second column and range from −39% to 39%. The distribution itself is computed by calculating the conditional probability of the net trade falling into a particular interval. The partition points for the net trade intervals are shown in the second row and range from −100% to 100%. The cells contain the conditional probabilities as percentages without the percentage sign. We provide the color coding of the cells to the right of this caption.

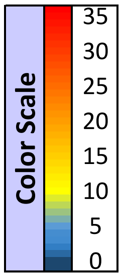

Figure 4-4: Probabilities of Security Liquidation and Large Buy Conditional on Past Returns

(a) Probability of Liquidation
Past Horizon: 20 days, Future Horizon: 60 days

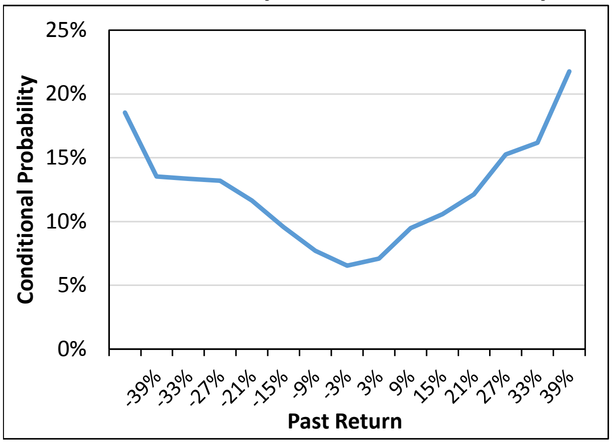

(b) Probability of Liquidation
Past Horizon: 20 days, Future Horizon: 5 days

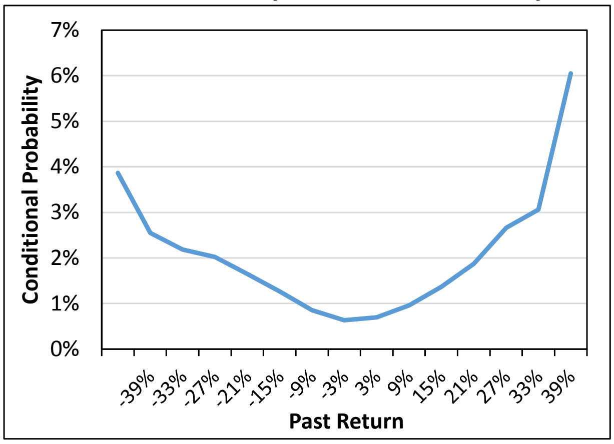

(c) Probability of Large Buy
Past Horizon: 20 days, Future Horizon: 60 days

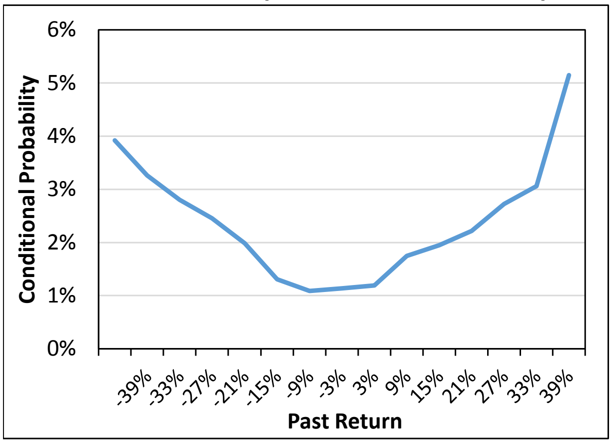

(d) Probability of Large Buy
Past Horizon: 20 days, Future Horizon: 5 days

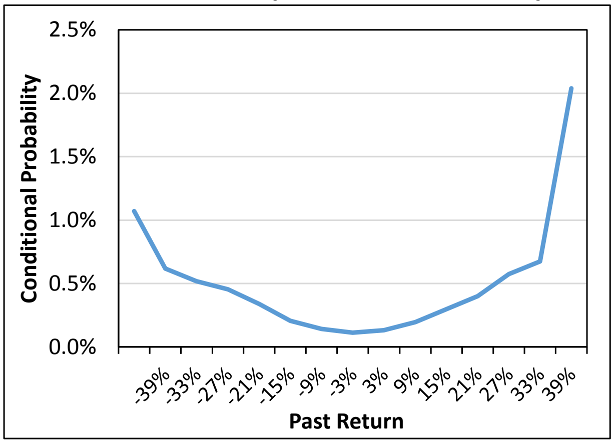

This figure plots the probability of liquidation and the probability of large buy conditional on returns over the past horizon. Analysis is done at the *security* level. Liquidation corresponds to a net trade of -100%, and a large buy corresponds to a net trade of 100% or more. We set the past horizon at 20 days and consider two cases for the future horizon: 60 days and 5 days.
All past return intervals have length 6%, except for the first interval that corresponds to a return of -39% or less and the last interval that corresponds to a return of over 39%. For each interval we calculate the conditional probability by dividing the number of observations corresponding to the particular net trade *and* past return falling into a particular interval by the total number of observations corresponding to the past return interval.
Standard errors do not exceed 0.08% for all conditional sell probabilities, and do not exceed 0.04% for all conditional buy probabilities. The confidence intervals are not plotted on the charts because they are very narrow.

Figure 4-5: Probabilities of Security Liquidation and Large Buy Conditional on Past Returns and Past Trading

(a) Probability of Liquidation, Trade vs. No Trade
Past Horizon: 20 days, Future Horizon: 60 days

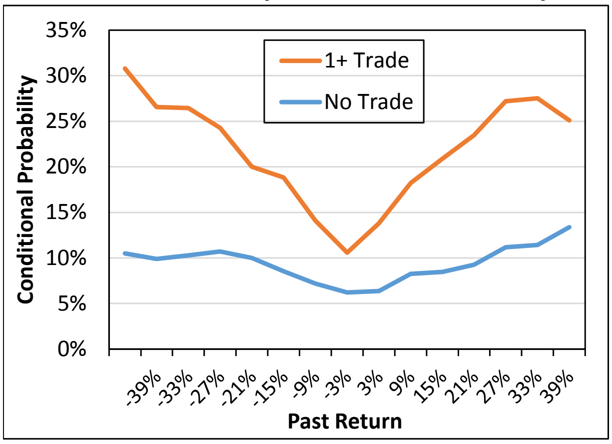

(b) Probability of Liquidation, Trade vs. No Trade
Past Horizon: 20 days, Future Horizon: 5 days

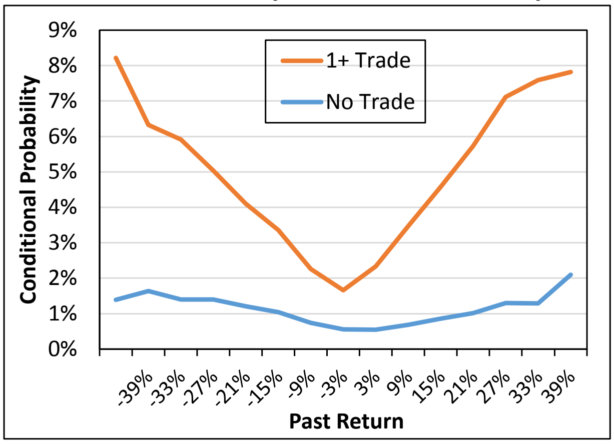

(c) Probability of Large Buy, Trade vs. No Trade
Past Horizon: 20 days, Future Horizon: 60 days

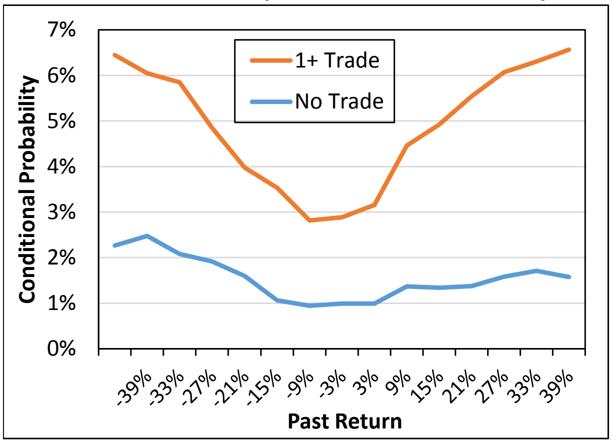

(d) Probability of Large Buy, Trade vs. No Trade
Past Horizon: 20 days, Future Horizon: 5 days

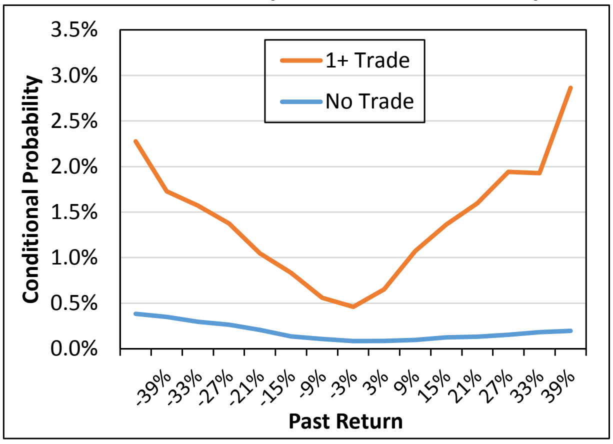

This figure plots the probability of liquidation and the probability of large buy conditional on returns and whether an investor traded over the past horizon. Analysis is done at the *security* level. Liquidation corresponds to a net trade of -100%, and a large buy corresponds to a net trade of 100% or more. We set the past horizon at 20 days and consider two cases for the future horizon: 60 days and 5 days. We split past trading behavior into two cases: when no trade was made over the past horizon and when at least one trade was made over the past horizon.
All past return intervals have length 6%, except for the first interval that corresponds to a return of -39% or less and the last interval that corresponds to a return of over 39%. For each interval we calculate the conditional probability by dividing the number of observations corresponding to the particular net trade *and* past return falling into a particular interval by the total number of observations corresponding to the past return interval.
Standard errors do not exceed 0.19% for all conditional sell probabilities, and do not exceed 0.10% for all conditional buy probabilities. The confidence intervals are not plotted on the charts because they are very narrow.

Figure 4-6: Probabilities of Portfolio Liquidation and Large Buy Conditional on Past Returns and Past Trading

(a) Probability of Liquidation, Trade vs. No Trade
Past Horizon: 20 days, Future Horizon: 60 days

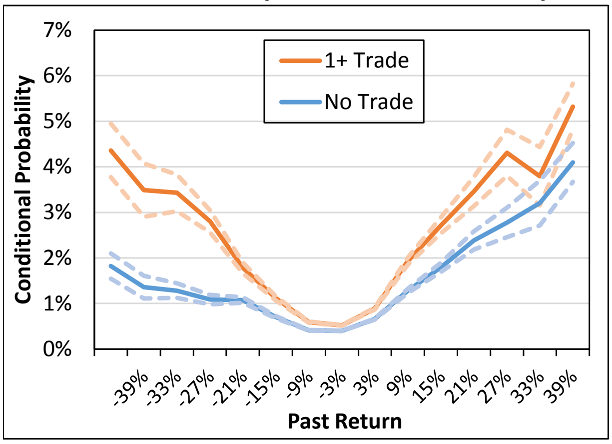

(b) Probability of Liquidation, Trade vs. No Trade
Past Horizon: 20 days, Future Horizon: 5 days

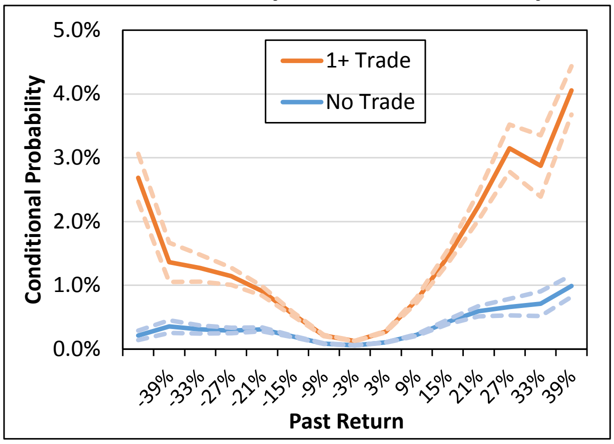

(c) Probability of Large Buy, Trade vs. No Trade
Past Horizon: 20 days, Future Horizon: 60 days

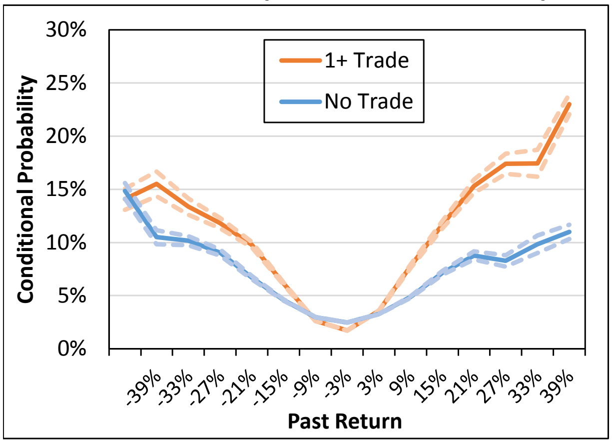

(d) Probability of Large Buy, Trade vs. No Trade
Past Horizon: 20 days, Future Horizon: 5 days

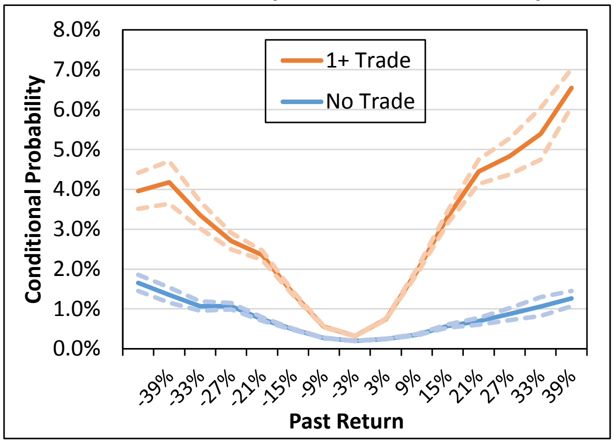

This figure plots the probability of liquidation and the probability of large buy conditional on returns and whether an investor traded over the past horizon. Analysis is done at the *portfolio* level. Liquidation corresponds to a net trade of -100%, and a large buy corresponds to a net trade of 100% or more. We set the past horizon at 20 days and consider two cases for the future horizon: 60 days and 5 days. We split past trading behavior into two cases: when no trade was made over the past horizon and when at least one trade was made over the past horizon.
All past return intervals have length 6%, except for the first interval that corresponds to a return of -39% or less and the last interval that corresponds to a return of over 39%. For each interval we calculate the conditional probability by dividing the number of observations corresponding to the particular net trade *and* past return falling into a particular interval by the total number of observations corresponding to the past return interval.
We plot the upper and lower bounds for the probability confidence intervals at the ± 2 standard error level; these are shown in light dashed lines.

Table 4.4: Steepness of Liquidation Prob. V-shape for Losses vs. Flat Returns

(a) Security Liquidations

Liquidation Losses / Flat		Future Horizon (days)											
		5	10	15	20	25	30	35	40	45	50	55	60
Past Horizon (days)	10	6.79	5.33	4.89	4.53	4.20	3.93	3.72	3.51	3.37	3.24	3.11	3.04
	20	6.11	5.11	4.63	4.29	4.01	3.76	3.49	3.30	3.17	3.06	2.95	2.83
	30	5.88	4.83	4.24	3.87	3.64	3.41	3.17	3.05	2.93	2.80	2.67	2.58
	40	5.39	4.63	4.17	3.82	3.60	3.40	3.18	3.02	2.89	2.76	2.65	2.56
	50	5.09	4.25	3.88	3.60	3.39	3.21	3.02	2.86	2.72	2.58	2.48	2.42
	60	4.61	3.99	3.73	3.48	3.28	3.08	2.88	2.75	2.63	2.52	2.42	2.38
	70	4.55	3.92	3.60	3.37	3.17	2.95	2.75	2.60	2.48	2.40	2.34	2.35
	80	4.39	3.86	3.53	3.23	3.00	2.79	2.62	2.48	2.38	2.30	2.25	2.26
	90	4.07	3.56	3.27	3.02	2.85	2.69	2.51	2.38	2.30	2.22	2.13	2.11
	100	4.05	3.51	3.30	3.10	2.91	2.73	2.54	2.38	2.27	2.19	2.11	2.09
	110	4.21	3.69	3.45	3.20	2.96	2.77	2.55	2.37	2.29	2.19	2.10	2.11
	120	4.39	3.75	3.39	3.15	2.89	2.67	2.50	2.30	2.22	2.14	2.09	2.05

(b) Portfolio Liquidations

Liquidation Losses / Flat		Future Horizon (days)											
		5	10	15	20	25	30	35	40	45	50	55	60
Past Horizon (days)	10	11.22	9.20	8.84	8.98	8.20	8.85	9.25	7.95	7.79	8.65	8.19	7.68
	20	12.93	10.47	9.50	9.32	8.40	7.79	6.83	6.29	6.57	6.54	6.56	6.18
	30	11.04	9.08	8.59	7.71	6.92	6.70	5.99	6.02	6.21	6.31	6.20	5.72
	40	10.49	8.98	7.53	6.90	6.44	6.04	5.84	5.81	5.78	5.44	5.07	4.85
	50	10.08	8.32	7.41	6.78	6.31	6.20	6.00	5.65	5.33	4.94	4.50	4.36
	60	9.96	7.94	7.35	6.87	6.45	6.03	5.64	5.26	4.79	4.36	4.01	3.89
	70	10.36	8.41	7.73	6.98	6.28	5.70	5.18	4.74	4.37	4.00	3.72	3.64
	80	9.87	8.26	7.40	6.52	5.84	5.30	4.74	4.45	4.02	3.68	3.56	3.59
	90	9.16	7.70	6.89	6.18	5.50	4.90	4.53	4.19	3.93	3.76	3.71	3.77
	100	7.93	6.79	6.22	5.67	5.11	4.60	4.27	4.10	3.94	3.82	3.86	4.01
	110	7.53	6.33	5.71	5.26	4.71	4.34	4.13	4.07	3.92	3.98	4.12	4.38
	120	6.52	5.60	5.04	4.84	4.62	4.30	4.17	4.22	4.20	4.16	4.28	4.52

These tables list the "steepness" of the V-shaped liquidation probability curve calculated as the conditional probability of liquidation following losses of -39% or more divided by the probability of liquidation following returns between between -3% and 3%. Analysis is done at the *security* level for the top table and at the *portfolio* level for the bottom table. We consider past horizons ranging from 10 to 120 days, and future horizons ranging from 5 to 60 days.
The cells in the table are color-coded according to their values, with green corresponding to higher values and red corresponding to lower values.

4.6 Conclusion and Next Steps

In conclusion, we have developed initial intuition that that will help us build models of trading behavior later in the project. We analyze the distribution of net trades and find that is it very distinct at the security and at the portfolio level. When trading securities, investors very frequently liquidate their positions, as well as often shift their positions by ±50% or integer multiples. For portfolios, most trades make up 0% to 10% of portfolio notional; we also observe substantial buying behavior. One of the reasons for this drastic difference in distributions could be portfolio "reshuffling", which constitutes liquidations of some securities and buying of other securities instead. The most important implication is that we should be considering two different models for security and for portfolio trading. If we aim to construct one "generic" model, then we need to incorporate factors from both the portfolio and the security level paradigms.

We have also looked at the probabilities of liquidation and large buys conditional on past returns. Investors tend to carry out more selling and buying following more extreme returns on their holdings, leading to "V-shaped" conditional probabilities for large net trades as a function of past returns. These reactions are more pronounced in situations where at least one trade was carried out over the past horizon, in comparison to when no trades were made. This behavior is consistent across different choices of past and future horizons, with shorter horizons leading to steeper V-shaped probability curves.

Of course, to get more robust conclusions along the lines of the above results, we would need to carry out a regression analysis, or potentially a classification scheme, such as CART or logistic regression. This approach would also allow us to account for more types of past trading behavior, so that we would not just use an indicator for if a trade was made in the past or not, but several variables describing past trading, including the number of trades, the trade dispersion in terms of sign, the timing of the trades, and how much of the investor portfolio they made up. We would also want to incorporate more features in the model, such as general market movements, investor demographics, and other variables describing past security and portfolio returns behavior at various horizons. Classification of investors into distinct types is also an important task, since it may be appropriate to develop separate models for different investor classes. Finally, we aim to apply machine-learning methods to this large dataset to hopefully obtain even better predictions of future investor behavior.

4.7 Appendix

4.7.1 Return and Net Trade Intervals

As discussed in the methodology, we calculate distributions by assigning net trades and past returns to disjoint intervals. For past returns we construct 15 intervals:

$$(-\infty, -39\%], (-39\%, -33\%], (-33\%, -27\%], \ldots, (-3\%, 3\%], \ldots, (33\%, 39\%], (39\%, \infty)$$

The choice of the bounds -39% and 39% is rather arbitrary but is done to capture the majority of possible past price moves. We ensure all intervals except the first and the last have the same interval length, so that it is easier to compare conditional probabilities across different buckets.

For the unconditional distribution of net trades we use 78 intervals:

$$(-\infty, -100\%), [-100\%, -100\%], (-100\%, -92\%], (-92\%, -84\%], \ldots,$$
$$(-12\%, -4\%], (-4\%, 4\%], (4\%, 12\%], \ldots, (492\%, 500\%], (500\%, \infty)$$

All intervals have length 8% except the following. The left-most interval of $(-\infty, -100\%)$ corresponds to shorting, or very extreme selling in the case of the dollar-based net trade measure. Since we exclude shorting from the sample (it is a very rare occurrence in the retail dataset), then the only way to get a net trade value of -100% or less is if an investor liquidated their position/portfolio; in that case the value is exactly -100%. For the dollar-based measure it is possible to get values less than -100%. We separate out the $[-100\%, 100\%]$ interval since it corresponds to liquidations. The right-most interval of $(500\%, \infty)$ corresponds to increasing the position(s) by a factor of five or more, and while the choice of the 500% bound is rather arbitrary, we believe that it allows us to capture almost all of the trading behavior, while still being able to investigate the right tail in detail.

For the conditional distributions we only use the share-based net trade measure, and therefore do not need to worry about values less than -100%. We also do not consider the right tail in as much detail as for the unconditional distribution, and therefore set the upper bound at 100%. Finally, since trades may happen at round numbers, we make sure that our intervals are symmetric around 0%. Based on this intuition, we use the following intervals:

$$[-100\%, -100\%], (-100\%, -90\%], (-90\%, -80\%], \ldots,$$
$$(-10\%, 0\%), [0\%, 0\%], (0\%, 10\%), [10\%, 20\%), \ldots, [90\%, 100\%), [100\%, \infty)$$

We separate out the $[-100\%, -100\%]$ interval since it corresponds to liquidations. We also separate out the $[0\%, 0\%]$ interval since it corresponds to no change in position. Note that we do not need to separate out this interval when calculating the unconditional distribution, since for that calculation, we restrict our observations to cases where an investor changed their position and therefore would never get 0% for the net trade.

Figure 4-7: Net Trade Distribution at Portfolio Level
Dollar-Based Approach

This figure plots the unconditional distribution of the net trade over the future horizon. Analysis is done at the *portfolio* level. We consider the 1-day and the 60-day future horizon. The net trade is calculated as the sum of (signed) dollar amounts traded, discounted by the time-weighted portfolio returns, as in equation (7).
The top two charts show the distribution over the whole interval $(-\infty, \infty)$, whereas the bottom two show the right tail distribution over the interval $(100\%, \infty)$. For each chart the distribution is computed by calculating the probability of the net trade falling into a particular interval. The endpoints of the intervals are shown on the $x-$axis.
All intervals have length of 8% except the following. For the top charts, the left-most interval corresponds to a net trade of less than -100%, the second left-most interval corresponds to a net trade exactly equal to -100%, while the right-most interval corresponds to a net trade of greater than 100%. For the bottom charts the right-most interval corresponds to a net trade of greater than 500%. Standard errors for the empirical probability corresponding to each interval are at most 0.02% (except for the intervals between -12% and 12%, where the bound is 0.05%) and are not shown in the charts.

Figure 4-8: Security Trade Distribution Conditional on Past Returns and Trading

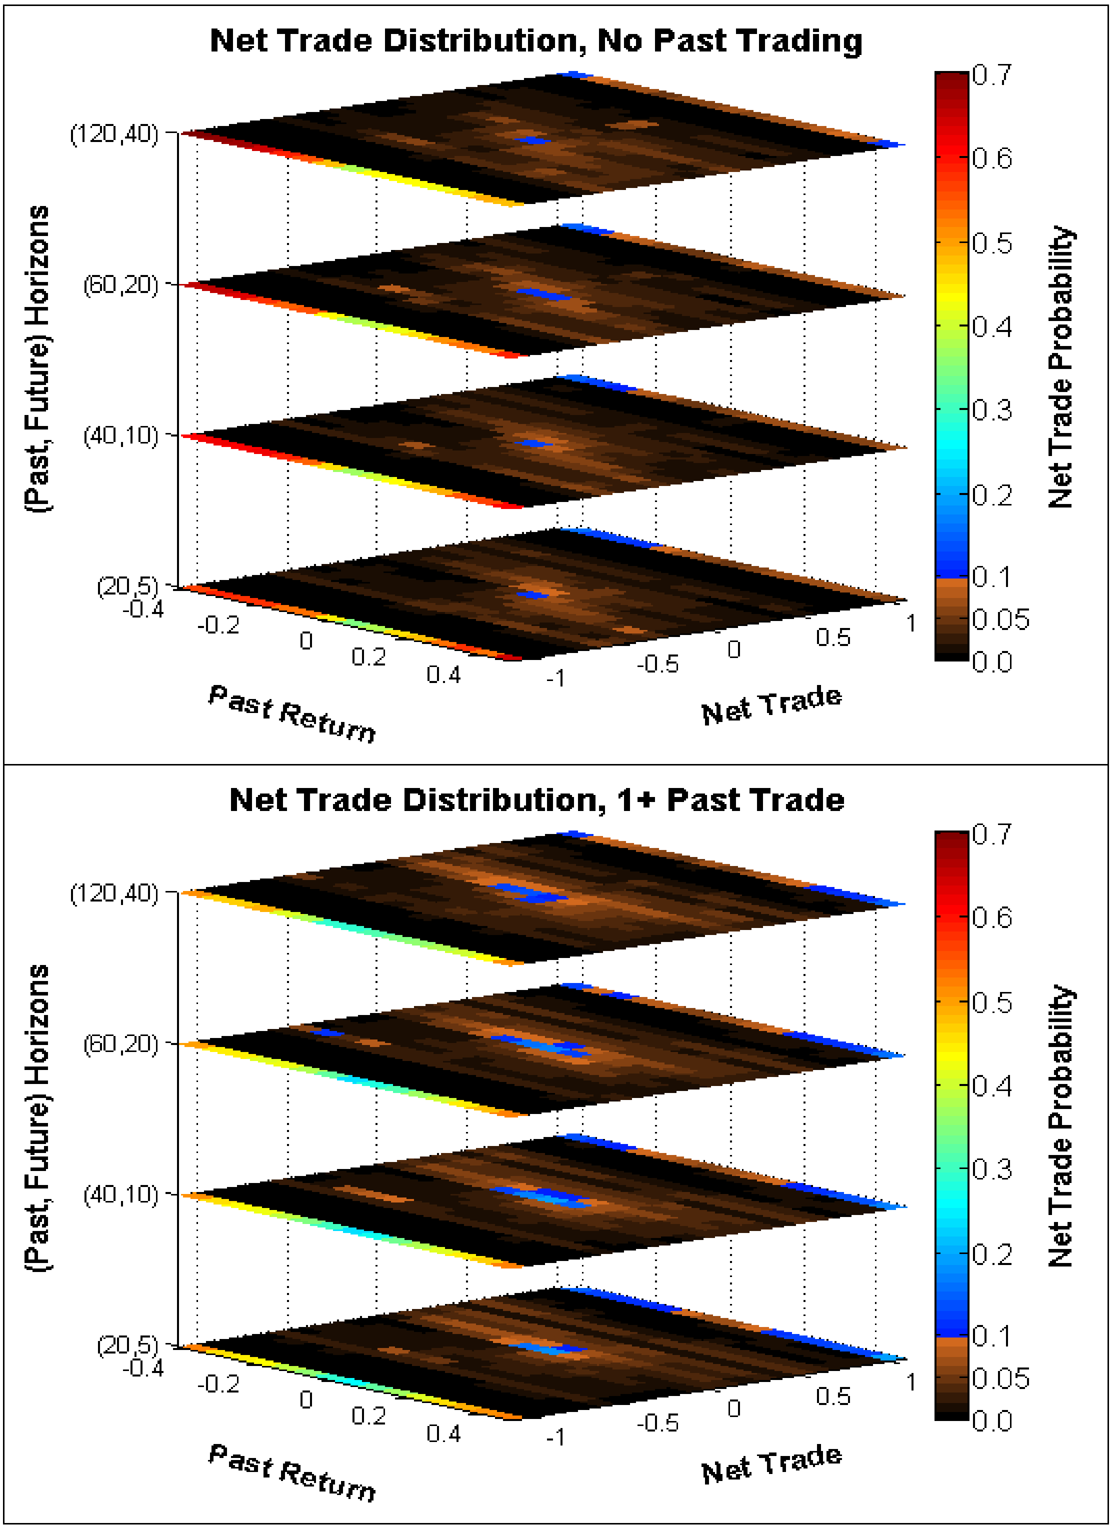

This figure plots the distribution of the net trade at the *security* level conditional on past returns and whether or not a trade was made over the past horizon. We consider four different cases for the past and future horizons. The heat map corresponding to each horizon case is constructed in the same way as in Table 4.2. We restrict our observations to those where the net trade is non-zero.

Figure 4-9: Portfolio Trade Distribution Cond. on Past Returns and Trading

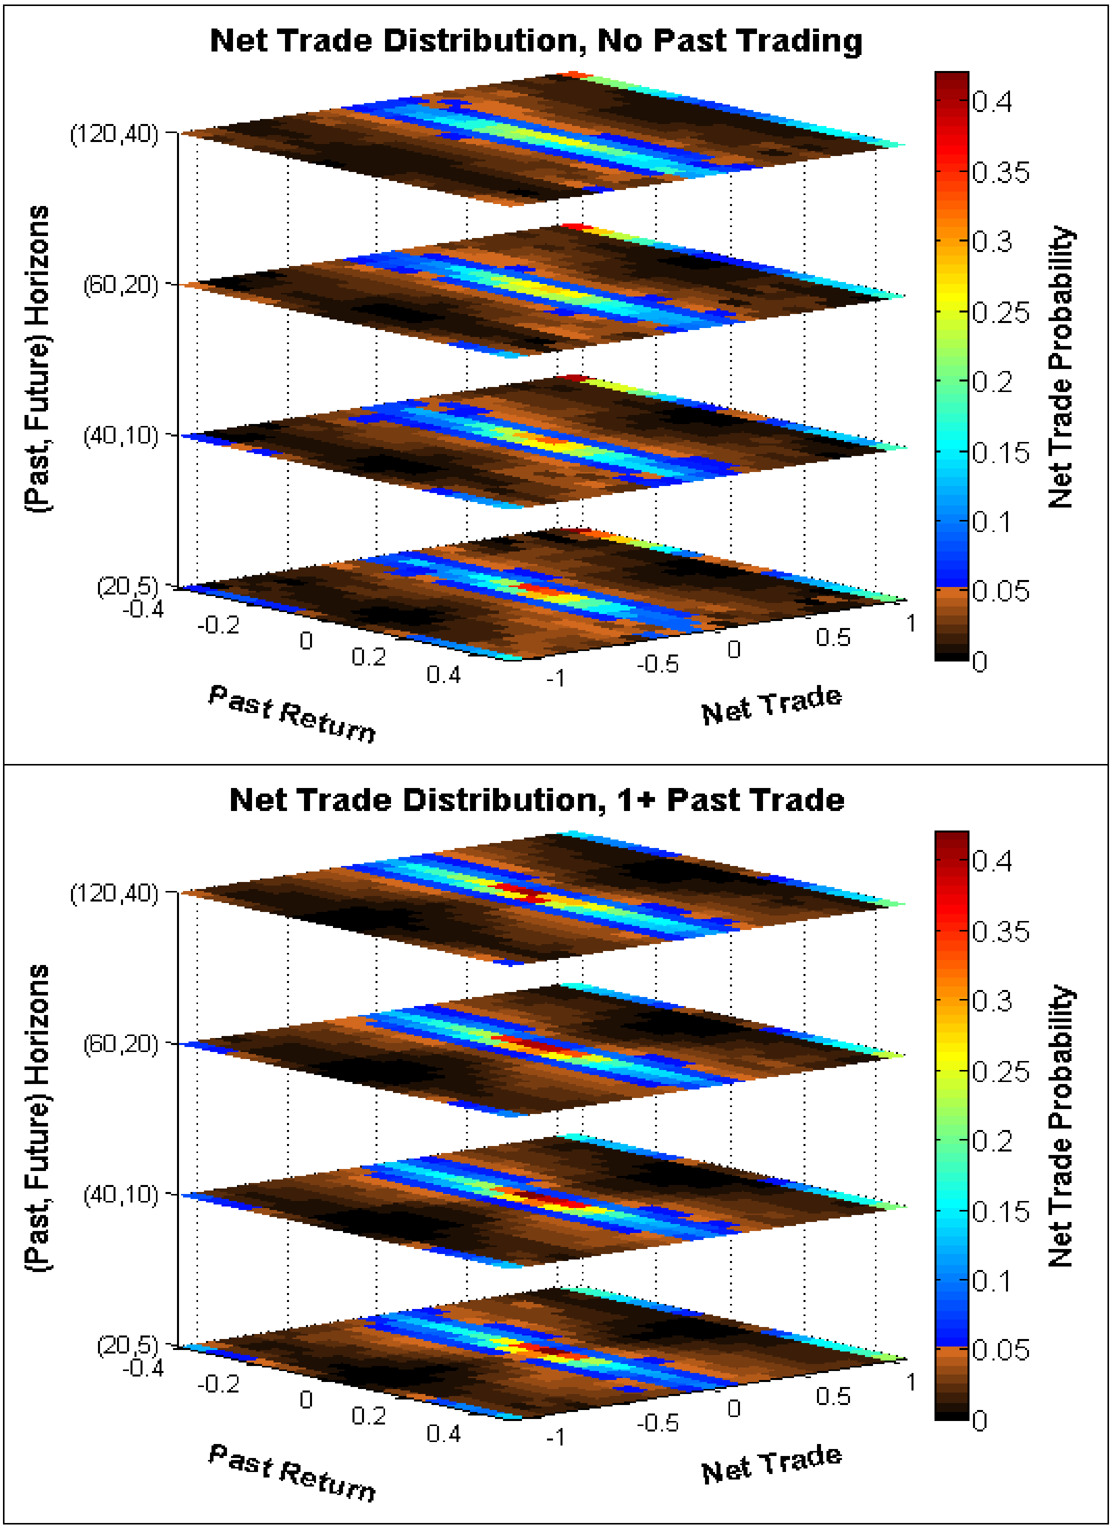

This figure plots the distribution of the net trade at the *portfolio* level conditional on past returns and whether or not a trade was made over the past horizon. We consider four different cases for the past and future horizons. The heat map corresponding to each horizon case is constructed in the same way as in Table 4.3. We restrict our observations to those where the net trade is non-zero.

Figure 4-10: Probabilities of Portfolio Liquidation and Large Buy Conditional on Past Returns

(a) Probability of Liquidation
Past Horizon: 20 days, Future Horizon: 60 days

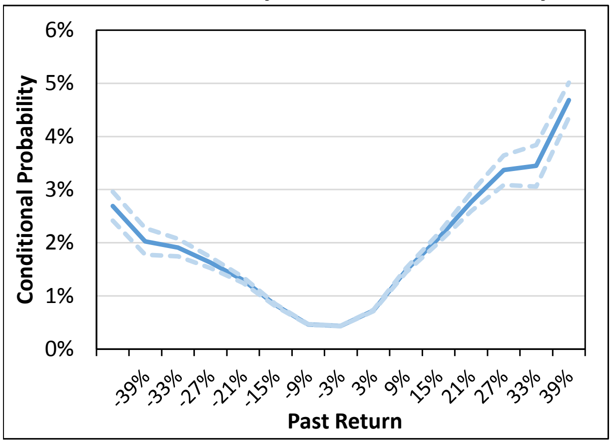

(b) Probability of Liquidation
Past Horizon: 20 days, Future Horizon: 5 days

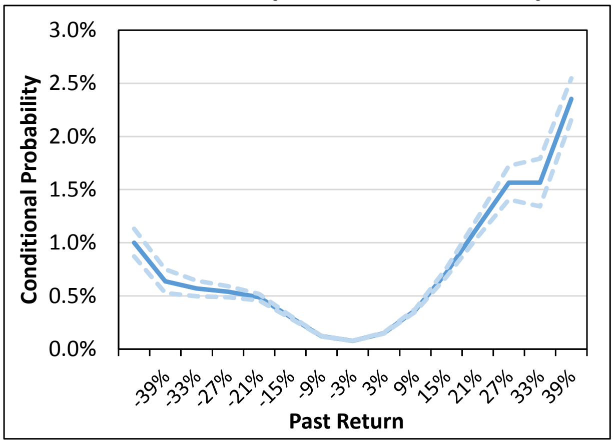

(c) Probability of Large Buy
Past Horizon: 20 days, Future Horizon: 60 days

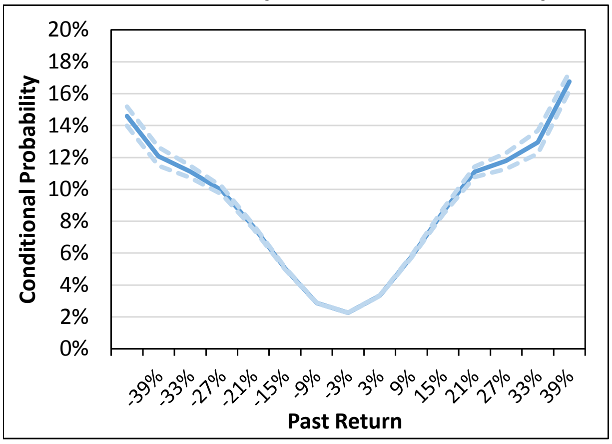

(d) Probability of Large Buy
Past Horizon: 20 days, Future Horizon: 5 days

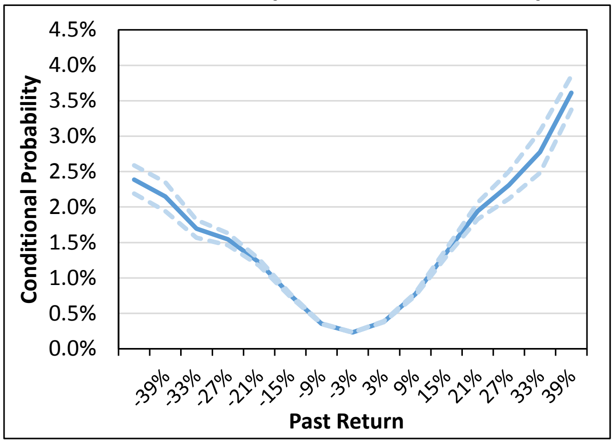

This figure plots the probability of liquidation and the probability of large buy conditional on returns over the past horizon. Analysis is done at the *portfolio* level. Liquidation corresponds to a net trade of -100%, and a large buy corresponds to a net trade of 100% or more. We set the past horizon at 20 days and consider two cases for the future horizon: 60 days and 5 days. All past return intervals have length 6%, except for the first interval that corresponds to a return of -39% or less and the last interval that corresponds to a return of over 39%. For each interval we calculate the conditional probability by dividing the number of observations corresponding to the particular net trade *and* past return falling into a particular interval by the total number of observations corresponding to the past return interval.
We plot the upper and lower bounds for the probability confidence intervals at the ± 2 standard error level; these are shown in light dashed lines.

Bibliography

V. Abramov, M. K. Khan, and R. A. Khan. A probabilistic analysis of the trading the line strategy. *Quantitative Finance*, 8(5):499–512, 2008.

Julie Agnew, Pierluigi Balduzzi, and Annika Sunden. Portfolio choice and trading in a large 401 (k) plan. *American Economic Review*, pages 193–215, 2003.

Gordon J Alexander and Mark A Peterson. An analysis of trade-size clustering and its relation to stealth trading. *Journal of Financial Economics*, 84(2):435–471, 2007.

Mark D Alicke, Mary L Klotz, David L Breitenbecher, Tricia J Yurak, and Debbie S Vredenburg. Personal contact, individuation, and the better-than-average effect. *Journal of personality and social psychology*, 68(5):804, 1995.

Brad M Barber and Terrance Odean. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *The journal of Finance*, 55(2):773–806, 2000.

Brad M Barber and Terrance Odean. Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly journal of Economics*, pages 261–292, 2001.

Brad M Barber and Terrance Odean. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies*, 21(2):785–818, 2008.

Nicholas Barberis, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. Extrapolation and bubbles. *Harvard and Yale unpublished manuscript*, 2015.

Nicholas Barberis, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. X-capm: An extrapolative capital asset pricing model. *Journal of Financial Economics*, 115(1):1–24, 2015.

Nicholas Barberis and Wei Xiong. What drives the disposition effect? an analysis of a long-standing preference-based explanation. *Journal of Finance*, 64(2):751–784, 2009.

Nicholas Barberis and Wei Xiong. Realization utility. *Journal of Financial Economics*, 104(2):251–271, 2012.

Manel Baucells, Martin Weber, and Frank Welfens. Reference-point formation and updating. *Management Science*, 57(3):506–519, 2011.

Itzhak Ben-David and David Hirshleifer. Are investors really reluctant to realize their losses? trading responses to past returns and the disposition effect. *Review of Financial Studies*, 25(8):2485–2532, 2012.

Shlomo Benartzi. Excessive extrapolation and the allocation of 401 (k) accounts to company stock. *Journal of Finance*, pages 1747–1764, 2001.

Georgina Benou and Nivine Richie. The reversal of large stock price declines: The case of large firms. *Journal of Economics and Finance*, 27(1):19–38, 2003.

Eta S Berner and Mark L Graber. Overconfidence as a cause of diagnostic error in medicine. *The American Journal of Medicine*, 121(5):S2–S23, 2008.

Michael W Brandt, Amit Goyal, Pedro Santa-Clara, and Jonathan R Stroud. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Review of Financial Studies*, 18(3):831–873, 2005.

Marc Bremer and Richard J. Sweeney. The reversal of large stock-price decreases. *The Journal of Finance*, 46(2):747–754, 1991.

Roger Buehler, Dale Griffin, and Michael Ross. Exploring the" planning fallacy": Why people underestimate their task completion times. *Journal of Personality and Social Psychology*, 67(3):366, 1994.

Bruce D Burns. When it is adaptive to follow streaks: Variability and stocks. In *Proceedings of the Twenty-fifth Annual Meeting of the Cognitive Science Society. Mahwah, NJ: Lawrence Erlbaum Associates*, 2003.

Colin Camerer and Dan Lovallo. Overconfidence and excess entry: An experimental approach. *American economic review*, pages 306–318, 1999.

Gavin Cassar. Are individuals entering self-employment overly optimistic? an empirical test of plans and projections on nascent entrepreneur expectations. *Strategic Management Journal*, 31(8):822–840, 2010.

James J Choi, David Laibson, Brigitte C Madrian, and Andrew Metrick. Reinforcement learning and savings behavior. *Journal of Finance*, 64(6):2515–2534, 2009.

Wen-I Chuang and Rauli Susmel. Who is the more overconfident trader? individual vs. institutional investors. *Journal of Banking & Finance*, 35(7):1626–1644, 2011.

Charles T Clotfelter and Philip J Cook. Notes: The “gambler’s fallacyâĂİ in lottery play. *Management Science*, 39(12):1521–1525, 1993.

Rachel Croson and James Sundali. The gamblerâĂŹs fallacy and the hot hand: Empirical data from casinos. *Journal of Risk and Uncertainty*, 30(3):195–209, 2005.

Magnus Dahlquist, José Vicente Martinez, and Paul Söderlind. Individual investor activity and performance. Technical Report 14-08, Swedish House of Finance Research Paper, 2014.

Kent Daniel, David Hirshleifer, and Avanidhar Subrahmanyam. Investor psychology and security market under-and overreactions. *Journal of Finance*, 53(6):1839–1885, 1998.

Chetan Dave, Catherine C Eckel, Cathleen A Johnson, and Christian Rojas. Eliciting risk preferences: When is simple better? *Journal of Risk and Uncertainty*, 41(3):219–243, 2010.

Werner PM De Bondt. Betting on trends: Intuitive forecasts of financial risk and return. *International Journal of Forecasting*, 9(3):355–371, 1993.

Leo De Haan and Jan Kakes. Momentum or contrarian investment strategies: evidence from dutch institutional investors. *Journal of Banking & Finance*, 35(9):2245–2251, 2011.

Richard Deaves, Erik Lüders, and Guo Ying Luo. An experimental test of the impact of overconfidence and gender on trading activity. *Review of finance*, page rfn023, 2008.

Kira Detko, Wilson Ma, and Guy Morita. Re-examining the hidden costs of the stop-loss. *University of Washington Working Paper*, 2008.

Daniel Dorn and Gur Huberman. Talk and action: What individual investors say and what they do. *Review of Finance*, 9(4):437–481, 2005.

Erik Ekstrom, Carl Lindberg, and Johan Tysk. Optimal liquidation of a pairs trade. *Advanced Mathematical Methods for Finance*, 2011.

Robin Erdestam and Olof Stangenberg. Efficiency of stop-loss rules. *University essay from Handelshögskolan i Stockholm/Institutionen för finansiell ekonomi*, 2008.

Sergei Esipov and Igor Vaysburd. On the profit and loss distribution of dynamic hedging strategies. *International Journal of Theoretical and Applied Finance*, 2(2):131–152, 1999.

Lei Feng and Mark S Seasholes. Do investor sophistication and trading experience eliminate behavioral biases in financial markets? *Review of Finance*, 9(3):305–351, 2005.

Stephen P. Ferris, Robert A. Haugen, and Anil K. Makhija. Predicting contemporary volume with historic volume at differential price levels: Evidence supporting the disposition effect. *The Journal of Finance*, 43(3):677–697, 1988.

Baruch Fischhoff, Paul Slovic, and Sarah Lichtenstein. Knowing with certainty: The appropriateness of extreme confidence. *Journal of Experimental Psychology: Human perception and performance*, 3(4):552, 1977.

Suzanne O'Curry Fogel and Thomas Berry. The disposition effect and individual investor decisions: the roles of regret and counterfactual alternatives. *Journal of Behavioral Finance*, 7(2):107–116, 2006.

Andrea Frazzini. The disposition effect and underreaction to news. *Journal of Finance*, 61(4):2017–2046, 2006.

Ryan Garvey and Anthony Murphy. Are professional traders too slow to realize their losses? *Financial Analysts Journal*, pages 35–43, 2004.

Thomas J George and Chuan-Yang Hwang. The 52-week high and momentum investing. *Journal of Finance*, 59(5):2145–2176, 2004.

Simon Gervais and Terrance Odean. Learning to be overconfident. *Review of Financial studies*, 14(1):1–27, 2001.

Thomas Gilovich, Robert Vallone, and Amos Tversky. The hot hand in basketball: On the misperception of random sequences. *Cognitive psychology*, 17(3):295–314, 1985.

Markus Glaser and Martin Weber. Which past returns affect trading volume? *Journal of Financial Markets*, 12(1):1–31, 2009.

Peter W. Glynn and Donald L. Iglehart. Trading securities using trailing stops. *Management Science*, 41(6):1096–1106, 1995.

Robin Greenwood and Andrei Shleifer. Expectations of returns and expected returns. *Review of Financial Studies*, pages 715–746, 2014.

John M Griffin, Jeffrey H Harris, and Selim Topaloglu. The dynamics of institutional and individual trading. *Journal of Finance*, 58(6):2285–2320, 2003.

Mark Grinblatt and Matti Keloharju. The investment behavior and performance of various investor types: a study of finland's unique data set. *Journal of Financial Economics*, 55(1):43–67, 2000.

Mark Grinblatt and Matti Keloharju. How distance, language, and culture influence stockholdings and trades. *The Journal of Finance*, 56(3):1053–1073, 2001.

Mark Grinblatt and Matti Keloharju. What makes investors trade? *Journal of Finance*, 56(2):589–616, 2001.

Mark Grinblatt and Matti Keloharju. Sensation seeking, overconfidence, and trading activity. *Journal of Finance*, 64(2):549–578, 2009.

Michael S Haigh and John A List. Do professional traders exhibit myopic loss aversion? an experimental analysis. *Journal of Finance*, 60(1):523–534, 2005.

Joop Hartog, Ada Ferrer-i Carbonell, and Nicole Jonker. Linking measured risk aversion to individual characteristics. *Kyklos*, 55(1):3–26, 2002.

Samuel M Hartzmark. The worst, the best, ignoring all the rest: The rank effect and trading behavior. *Review of Financial Studies*, 28(4):1024–1059, 2015.

Joel Hasbrouck. Trading costs and returns for us equities: Estimating effective costs from daily data. *The Journal of Finance*, 64(3):1445–1477, 2009. `http://people.stern.nyu.edu/jhasbrou/Research/GibbsCurrent/gibbsCurrentIndex.html`.

David Hirshleifer. Investor psychology and asset pricing. *The Journal of Finance*, 56(4):1533–1597, 2001.

Charles A Holt and Susan K Laury. Risk aversion and incentive effects. *American Economic Review*, 92(5):1644–1655, 2002.

Gerd Infanger. Dynamic asset allocation strategies using a stochastic dynamic programming approach. *Handbook of asset and liability management*, 1:199–251, 2006.

Andrew Jackson. The aggregate behaviour of individual investors. *Available at SSRN 536942*, 2003.

Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):65–91, 1993.

Joseph Johnson, Gerard J Tellis, and Deborah J MacInnis. Losers, winners, and biased trades. *Journal of Consumer Research*, 32(2):324–329, 2005.

Magne JøRgensen, Karl Halvor Teigen, and Kjetil MoløKken. Better sure than safe? over-confidence in judgement based software development effort prediction intervals. *Journal of Systems and Software*, 70(1):79–93, 2004.

Daniel Kahneman and Amos Tversky. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.

Kathryn M. Kaminski and Andrew W. Lo. When do stop-loss rules stop losses? *Journal of Financial Markets*, 18:234–254, 2014.

Ron Kaniel, Gideon Saar, and Sheridan Titman. Individual investor trading and stock returns. *Journal of Finance*, 63(1):273–310, 2008.

Joshua Klayman, Jack B Soll, Claudia González-Vallejo, and Sema Barlas. Overconfidence: It depends on how, what, and whom you ask. *Organizational behavior and human decision processes*, 79(3):216–247, 1999.

Philipp Koellinger, Maria Minniti, and Christian Schade. "i think i can, i think i can": Overconfidence and entrepreneurial behavior. *Journal of Economic Psychology*, 28(4):502–527, 2007.

Richard Krivo. Don't take profits too quickly or stay in losing trades too long. *DailyFX Trading Tips*, 2012. `http://www.dailyfx.com/forex/education/trading_tips/post_of_the_day/2012/05/31/Dont_Take_Profits_Too_Quickly_or_Stay_in_Losing_Trades_Too_Long.html`.

Adam Y.C. Lei and Huihua Li. The value of stop loss strategies. *Financial Services Review*, 18(1):23–51, 2009.

Chien-Huang Lin, Wen-Hsien Huang, and Marcel Zeelenberg. Multiple reference points in investor regret. *Journal of Economic Psychology*, 27(6):781–792, 2006.

Andrew W. Lo and Craig MacKinlay. Stock market prices do not follow random walks: Evidence from a simple specification test. *The Review of Financial Studies*, 1(1):41–66, 1988.

Robert Macrae. The hidden cost of the stoploss. *AIMA Journal, Arcus Investment*, 2005.

Ulrike Malmendier and Stefan Nagel. Depression babies: Do macroeconomic experiences affect risk-taking? Technical report, National Bureau of Economic Research, 2009.

Robert Masters. Study examines investors' risk-taking propensities. *Journal of Financial Planning*, 2(3), 1989.

Robert C. Merton. Lifetime portfolio selection under uncertainty; the continuous-time case. *Review of Economics and Statistics*, 51(3):247–257, 1969.

Ciamac C Moallemi and Mehmet Sağlam. Dynamic portfolio choice with linear rebalancing rules. *Available at SSRN 2011605*, 2015.

Pamela C Moulton. You canâĂŹt always get what you want: Trade-size clustering and quantity choice in liquidity. *Journal of Financial Economics*, 78(1):89–119, 2005.

Margaret A Neale and Max H Bazerman. The effects of framing and negotiator overconfidence on bargaining behaviors and outcomes. *Academy of Management Journal*, 28(1):34–49, 1985.

Clifford Nowell and Richard M Alston. I thought i got an a! overconfidence across the economics curriculum. *Journal of Economic Education*, 38(2):131–142, 2007.

Terrance Odean. Are investors reluctant to realize their losses? *Journal of Finance*, 53(5):1775–1798, 1998.

Terrance Odean. Do investors trade too much? *American Economic Review*, 89(5):1279–1298, 1999.

An T Oskarsson, Leaf Van Boven, Gary H McClelland, and Reid Hastie. What's next? judging sequences of binary events. *Psychological bulletin*, 135(2):262, 2009.

Anne-Marie Pålsson. Does the degree of relative risk aversion vary with household characteristics? *Journal of economic psychology*, 17(6):771–787, 1996.

Moustapha Pemy. Optimal selling rule in a regime switching lévy market. *International Journal of Mathematics and Mathematical Sciences*, 2011.

Marcelo Perlin. Ms_regress – the matlab package for Markov regime-switching models. *Available at SSRN 1714016*, 2012.

James M. Poterba and Lawrence H. Summers. Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics*, 22(1):27 – 59, 1988.

Matthew Rabin and Dimitri Vayanos. The gambler's and hot-hand fallacies: Theory and applications. *The Review of Economic Studies*, 77(2):730–778, 2010.

Daniel Richards, Janette Rutterford, and Mark Fenton-O'Creevy. Do stop losses work? the disposition effect, stop losses and investor demographics. *The Open University Business School Working Paper*, 2011.

Pavel G. Savor. Stock returns after major price shocks: The impact of information. *Journal of Financial Economics*, 106(3):635 – 659, 2012.

Zur Shapira and Itzhak Venezia. Patterns of behavior of professionally managed and independent investors. *Journal of Banking & Finance*, 25(8):1573–1587, 2001.

Hersh Shefrin. *Behavioralizing Finance*. Now Publishers Inc, 2010.

Hersh Shefrin and Meir Statman. The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance*, 40(3):777–790, 1985.

Bergsveinn Snorrason and Garib Yusupov. Performance of stop-loss rules vs buy-and-hold strategy. *Lund University School of Economics and Management*, 2009.

Fred H Speece Jr and John W Rogers Jr. Conversation with a value guru. *CFA Institute Conference Proceedings Quarterly*, 27(4):1–7, 2010.

Michal Ann Strahilevitz, Terrance Odean, and Brad M Barber. Once burned, twice shy: How naïve learning, counterfactuals, and regret affect the repurchase of stocks previously sold. *Journal of Marketing Research*, 48(SPL):S102–S120, 2011.

Ola Svenson. Are we all less risky and more skillful than our fellow drivers? *Acta Psychologica*, 47(2):143–148, 1981.

Richard H Thaler, Amos Tversky, Daniel Kahneman, and Alan Schwartz. The effect of myopia and loss aversion on risk taking: An experimental test. *The Quarterly Journal of Economics*, pages 647–661, 1997.

Martin Weber and Colin F Camerer. The disposition effect in securities trading: An experimental analysis. *Journal of Economic Behavior & Organization*, 33(2):167–184, 1998.

Martin Weber and Frank Welfens. The follow-on purchase and repurchase behavior of individual investors: An experimental investigation. *Die Betriebswirtschaft*, 71:139–154, 2011.

Danielle D Winchester, Sandra J Huston, and Michael S Finke. Investor prudence and the role of financial advice. *Journal of Financial Service Professionals*, 65(4), 2011.

Alan S. Wong, Bernardo J. Carducci, and Alan Jay White. Asset disposition effect: The impact of price patterns and selected personal characteristics. *Journal of Asset Management*, 7(3):291–300, 2006.

Kenton K Yee. Deep-value investing, fundamental risks, and the margin of safety. *Journal of Investing*, 17(3):35–46, 2008.

Hanqin Zhang and Qing Zhang. Trading a mean-reverting asset: Buy low and sell high. *Automatica*, 44(6):1511–1518, 2008.

Qing Zhang. Stock trading: An optimal selling rule. *SIAM Journal on Control and Optimization*, 40(1):64–87, 2001.